UNITED STATES DEPARTMENT OF THE INTERIOR

GEOLOGICAL SURVEY

A NEW EVENT DETECTOR DESIGNED FOR THE SEISMIC RESEARCH OBSERVATORIES

by

JAMES N. MURDOCK

and

CHARLES R. HUTT

OPEN-FILE REPORT 83-785

OCTOBER 1983

ALBUQUERQUE, NEW MEXICO

CONTENTS

Abstract

A new short-period event detector has been implemented on the Seismic Research Observatories. For each signal detected, a printed output gives estimates of the time of onset of the signal, direction of the first break, quality of onset, period and maximum amplitude of the signal, and an estimate of the variability of the background noise. On the SRO system, the new algorithm runs ∿2.5X faster than the former (power level) detector. This increase in speed is due to the design of the algorithm: all operations can be performed by simple shifts, additions, and comparisons (floating point operations are not required). Even though a narrow-band recursive filter is not used, the algorithm appears to detect events competitively with those algorithms that employ such filters. Tests at Albuquerque Seismological Laboratory on data supplied by Blandford suggest performance commensurate with the on-line detector of the Seismic Data Analysis Center, Alexandria, Virginia.

Key Words: Digital event detector, signal detector, signal analyzer, Seismic Research Observatories, SRO, ASRO

INTRODUCTION

The purpose of this report is to describe an event detector that has been implemented on the Seismic Research Observatories (SRO) and the Modified High-Gain Long-Period Observatories (ASRO). The new detector is an improvement over the former one that was described by Unitech (1974, pp. 42-59). The former detector, a power level algorithm, employed recursive filters. In some instances, these filters exhibited numerical overflow. Also, the former detector required a large percentage of the available processing time, due to software multiply and divide operations in the recursive filters.

The new detector operates on the digital data of the short-period vertical channel of each of the SRO and ASRO systems, as did the power detector. (The short-period response of the SRO was described by Peterson et al. (1976) and the short-period response of the ASRO is similar to the SRO response.) The output of the new digital detector is designed to emulate the analysis by a human of the short-period SRO helicorder display. For each signal detected, this printed output gives estimates of time of onset of the signal, direction of the first break, quality of the onset and first break determinations, period and maximum amplitude of the signal, and an estimate of the variability of the background noise.

The new algorithm has been designed for programming on very elementary computers: all operations, including filtering, can be implemented by simple additions, shifts, and comparisons. This design has eliminated the overflow problem of the former detector and has made available much more processing time for other operations.

OVERVIEW OF FLOW OF OPERATIONS

The overview of the flow of the operations is shown by Figure 1 and described below. The numerals below correspond to the notation on the flow chart (Figure 1).

I. The input time series is filtered. Although any type of filter could be employed, we use a simple sum and difference (finite impulse response, FIR) bandpass filter on the SRO and ASRO systems.

II. Relative maximums and minimums (peaks and troughs) of the filtered series are found, and successive peaks and troughs are differenced. These differenced values, together with their associated times, are the time series (called the P-T series) that is processed by the remainder of the detector. Each associated time is for the last sample of the subset that defined the P-T value; i.e. for a P-T value that defines a trough-to-peak amplitude, the P-T time is that of the peak, and for a peak-to-trough amplitude, the P-T time is that of the trough.

III. An estimate (s') of the dispersion (variability) of the P-T series is made.

IV. Three thresholds (Th1, Th2, Th3) are calculated. All three are a function of s', with Th1>Th2>Th3. Thresholds Th1 and Th2 are used to detect events, and TH3 is used to search for onsets of detected events.

V. Each new value in the P-T series is compared to Th1 and Th2 that (as described above) are related to the dispersion of the P-T series. A detection is declared if, in a fixed time window (normally 4 sec) one value of the rectified P-T series exceeds Th1 and two other rectified values exceed Th2, or if m rectified values exceed Th2 only (normally m = 4), subject to two restrictions: namely, additional processing by two more windows that perform quasi-bandpass filtering. These two windows, named Filhi and Fillo, are subwindows of the 4-second window, and they are described in the next section.

VI. When a detection has been made, Th3 is used to search for the time of onset of the signal. The search begins two P-T values before the first one that exceeded Th2.

VII. Parameters that describe the signal are estimated and output. These include the onset time, the direction of first break, quality of onset and first break, average period and maximum amplitude of the first 4 cycles of the signal, and s'.

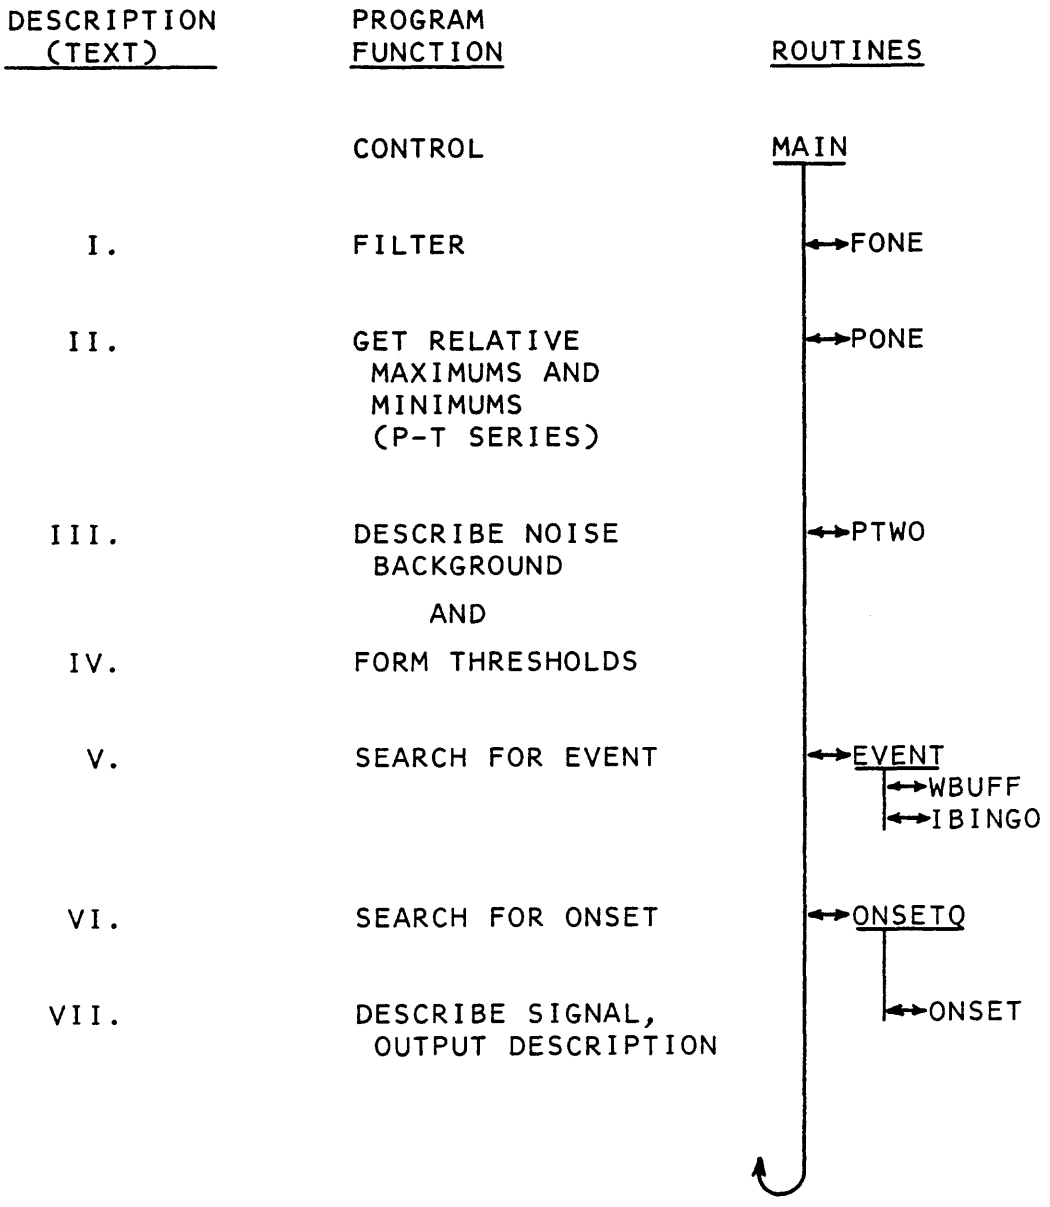

Figure 1. General flow chart of the algorithm. Roman numerals on the left of the illustration correspond to those describing the algorithm in the text. Routines called are shown on the right of the illustration.

DETAILS OF THE ALGORITHM

Whereas much of the flow of the program can be easily seen from the general flow chart (Figure 1) together with the general descriptions, some of the procedures need further explanation.

Filtering the Input Data

Both formal and informal filtering are performed. The formal filter currently used for the SRO is a simple sum and difference filter whose response is given in Figure 2. Also shown there is the attenuation of the first one-half cycle of a transient sine wave, relative to the steady-state response, as a function of frequency; in the band 0.5-2.5 Hz, the first break is attenuated a maximum of only 3.7 dB. In contrast a Butterworth 4-pole bandpass filter (0.5-2.5 Hz) attenuates the first break by a maximum of 8 dB in this band. If microseisms or cultural noise conditions require it, we are prepared to implement a narrow-band filter in the detector. However, to preserve information on the first break of the signal, we prefer the simple filter of Figure 2. Should experience dictate the need for a narrower band-pass filter, or one with a steeper roll-off, we may implement another FIR filter with a response as in Figure 3. An FIR filter with this response can be implemented with ten coefficients, all of which are integer powers of 2.

The algorithm employs four different types of informal filtering. The first is performed in PONE and the remainder are performed in EVENT. As described previously, the P-T series is formed from the output of the formal band-pass filter. Since the P-T series is made by differencing the successive maximums and minimums, forming it attenuates low frequency information. This type of informal filtering is exemplified by model data of Figure 4; it demonstrates how data in the pass band of the filter (2 Hz) might attenuate data of lower frequencies. Second, for a detection to occur, at least three values (not necessarily consecutive values) of the rectified P-T series must exceed Th2 in the moving time window (typically 4 sec). Thus, for a steady-state sine wave, signals having periods longer than 4.0 sec will not be detected, regardless of amplitude. Third, successive rectified P-T values greater than Th2 (in the 4-sec window) must occur within a subwindow whose length is determined by the input parameter Fillo (typically 2 sec). Otherwise, if the time between two successive values is greater than Fillo, the beginning of the 4-second window will be moved to the position of the last P-T value that exceeded Th2. Normally, Fillo is used to fine tune the detector. Fourth, when a rectified P-T value exceeds Th2, winnowing for a time determined by the parameter Filhi (typically 0.2 sec) is performed: P-T values in this small subwindow are not counted towards a detection, regardless of their amplitude. In some instances, winnowing is used to reject "ringing" of the filter. When used in this manner, the value for Filhi should be greater than the time difference between the first minimum and the first maximum of the impulse response of the filter (see Figure 5). The rejection of "ringing" is useful mainly when Th2 is set to a very small value (i.e. $\leq 1.0 \times s'$ where typical values are $\geq 1.5 \times s'$). Also, winnowing may be used to reject spikes and short bursts of cultural noise occasionally found in the input time series.

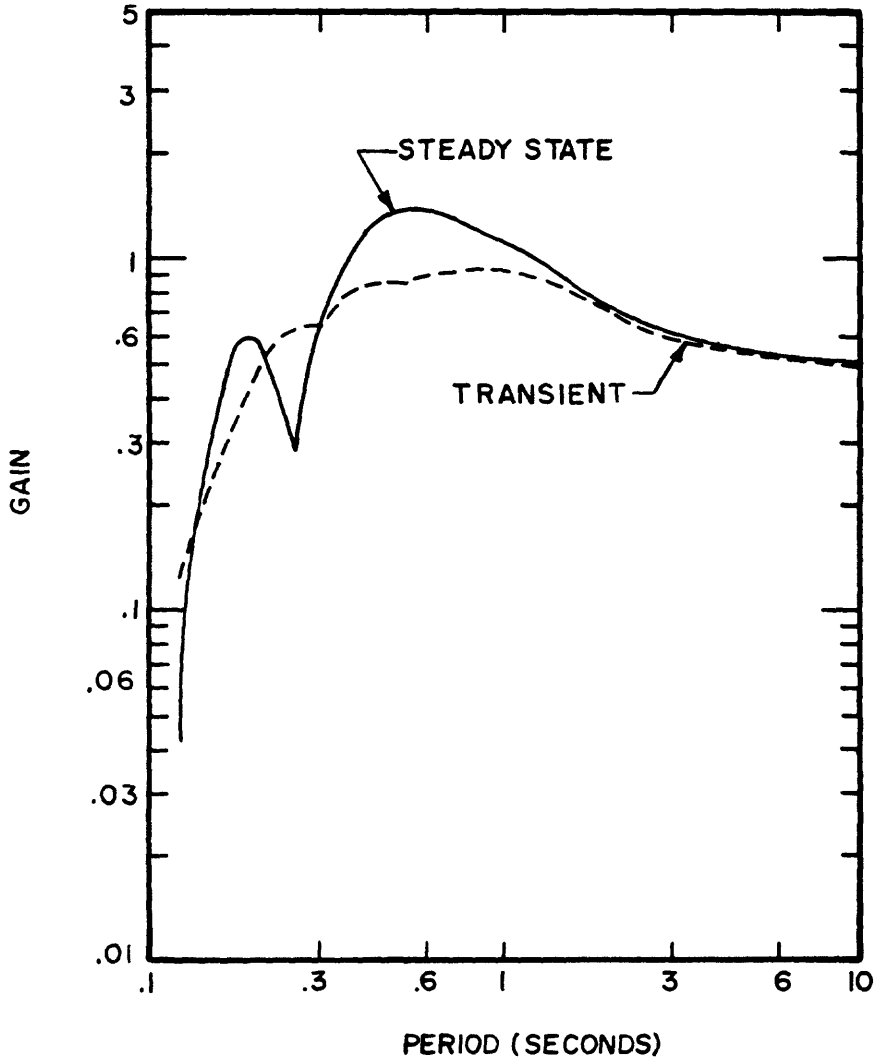

Filter 1 (F1)

$$g_m = \sum_{n=0}^{N} b_n f_{m-n}$$

$N=8$

f_m: input time series

g_m: output time series

b_n: coefficients

n	b_n
0	0
1	0.25
2	0.5
3	0.25
4	0
5	0
6	-0.125
7	-0.25
8	-0.125

Figure 2. Steady-state response of filter F1 together with its response to the first one-half cycle of a sine-wave transient. Filter equation and coefficients are shown.

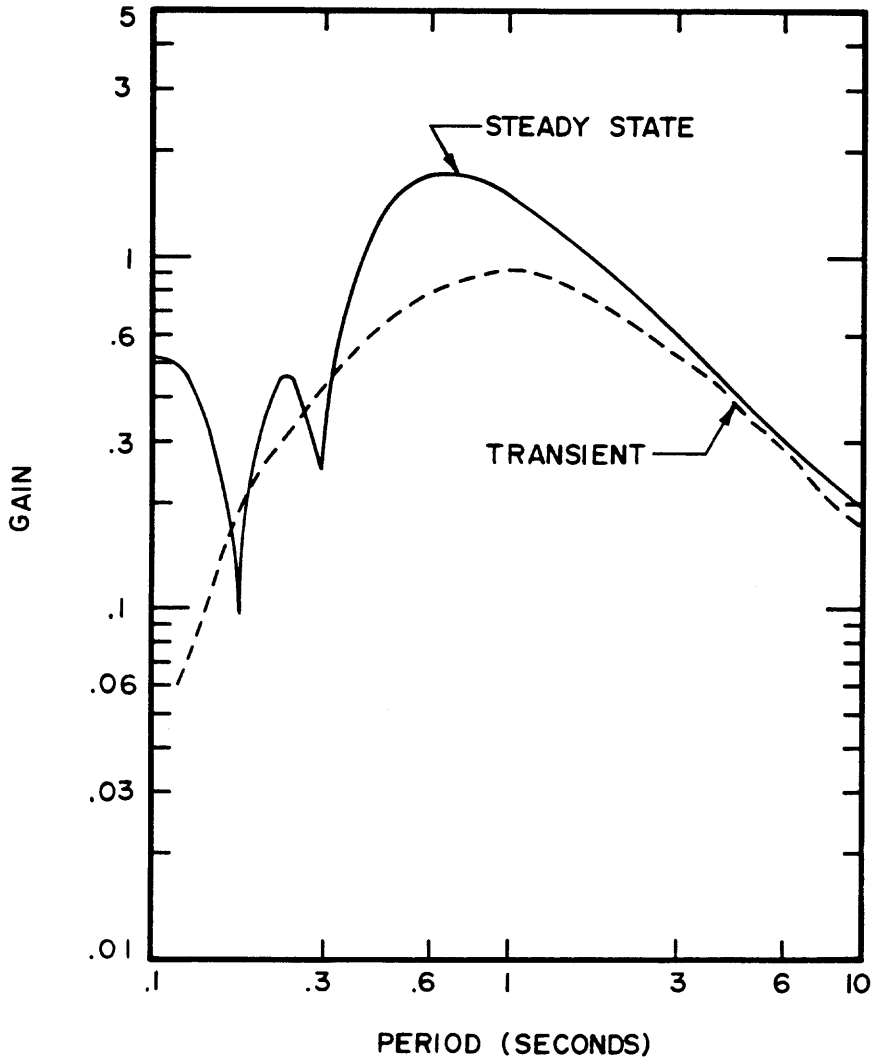

Filter 7 (F7)

$$g_m = \sum_{n=0}^{N} b_n f_{m-n}$$

N=10

f_m: input time series

g_m: output time series

b_n: coefficients

n	b_n
0	0
1	0.125
2	0.25
3	0.25
4	0.25
5	0.125
6	0
7	0
8	-0.5
9	-0.25
10	-0.25

Figure 3. Response of a potential replacement for F1. We prefer F1 because of its superior response to a transient. Filter equation and coefficients are shown.

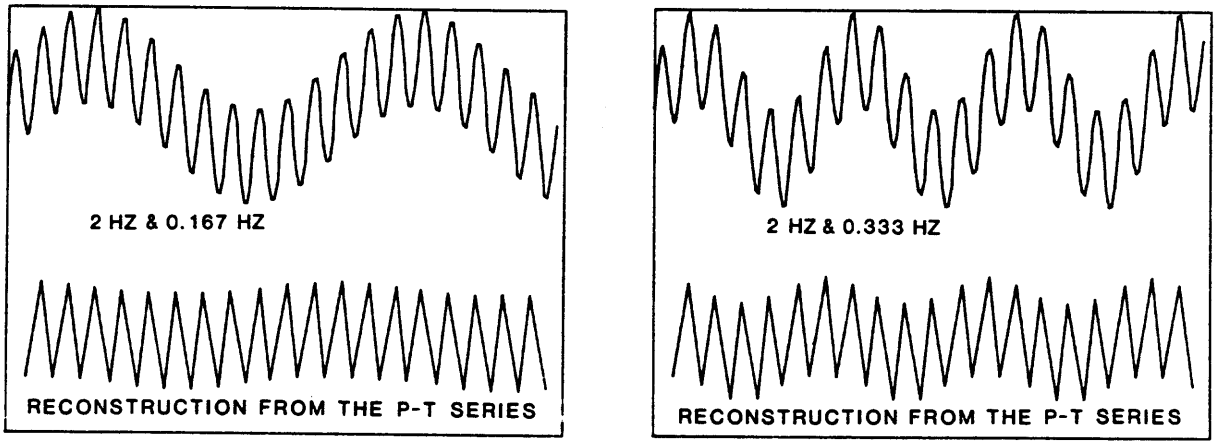

Figure 4. Illustration of informal filtering by differencing. When the P-T series is formed, low frequency information (i.e., much lower than the center frequency of the filter) is attenuated. The left illustration shows the effect of differencing when the input is the sum of a 2 Hz and a 0.167 Hz signal (both of the same amplitude); the right, the effect of differencing on the sum of a 2 Hz and a 0.333 Hz signal.

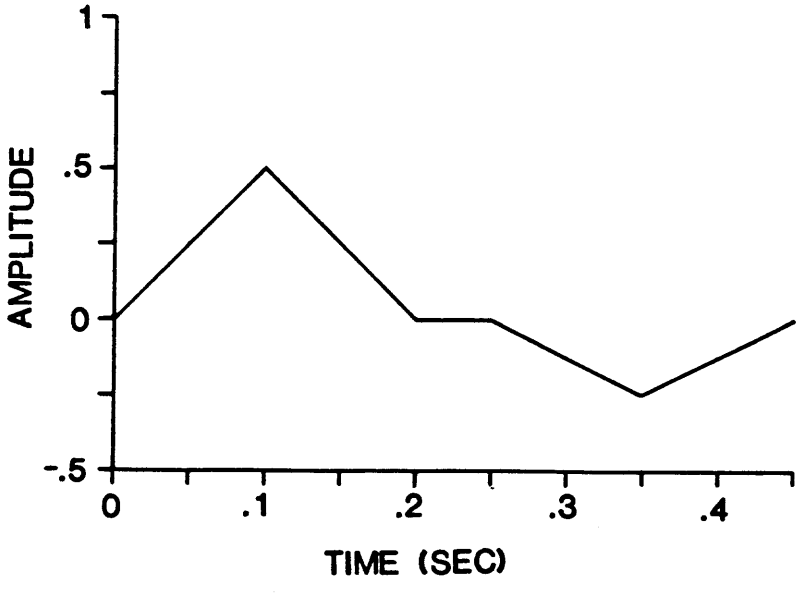

Figure 5. Impulse response of F1. False alarms can be reduced by deactivating detection (winnowing) long enough to allow passage of the wavelet that corresponds to the first minimum of the response (above). Winnowing is accomplished by using the input parameter Filhi. Filhi defines a short window length, not the corner frequency of a formal bandpass filter.

Estimating the Dispersion of the Background Noise and Forming the Thresholds

The dispersion (variability) of the background noise is estimated in two routines, PONE and PTWO. In PONE, the rectified P-T values are compared to a threshold for the noise, Thx (typically 1.5625×s'). Values less than Thx are put into a 20-element array, and when this array is full, PTWO is called. PTWO picks the maximum of the 20 rectified values, re-initializes the 20-element array, and averages the current maximum value with the 15 previous ones. The average value thus obtained is s'. For zero-mean normally distributed P-T values, s' is an estimate of twice the sample standard deviation of the P-T values. Examples of the correspondence of s' to s (the sample standard deviation of the P-T values) are given in Figures 6A-6J. Although there are exceptions (Figure 6A), typically the correspondence is approximately linear.

The two thresholds for detecting signals, Th1 and Th2, are formed from s' by using factors Xth1 and Xth2: Th1 = Xth1×s', Th2 = Xth2×s'. The threshold for timing the onset Th3 is formed in the same manner: Th3 = Xth3×s'. Values for Xth1, Xth2, and Xth3 are selected by the operator; typical values are Xth1 = 2., Xth2 = 1.5, Xth3 = 1.

Figure 6. The value s' plotted as a function of the sample standard deviation, s. To avoid running of ink on the plots, only every 16th point was plotted. Data for ten stations are shown on the next five pages.

		INTERVAL OF ESTIMATE, 1980	
FIGURE	STATION	DAY	HOUR:MINUTE
6A	ANMO	288	15:00
		289	15:00
6B	BCAO	288	06:00
		289	06:00
6C	CHTO	283	00:00
		284	00:00
6D	CTAO	279	12:00
		280	12:00
6E	GUMO	281	00:00
		282	00:00
6F	GRFO	282	14:00
		283	14:00
6G	KONO	275	17:00
		276	17:00
6H	MAJO	282	00:00
		283	00:00
6I	TATO	285	00:00
		286	00:00
6J	ZOBO	287	00:00
		288	00:00

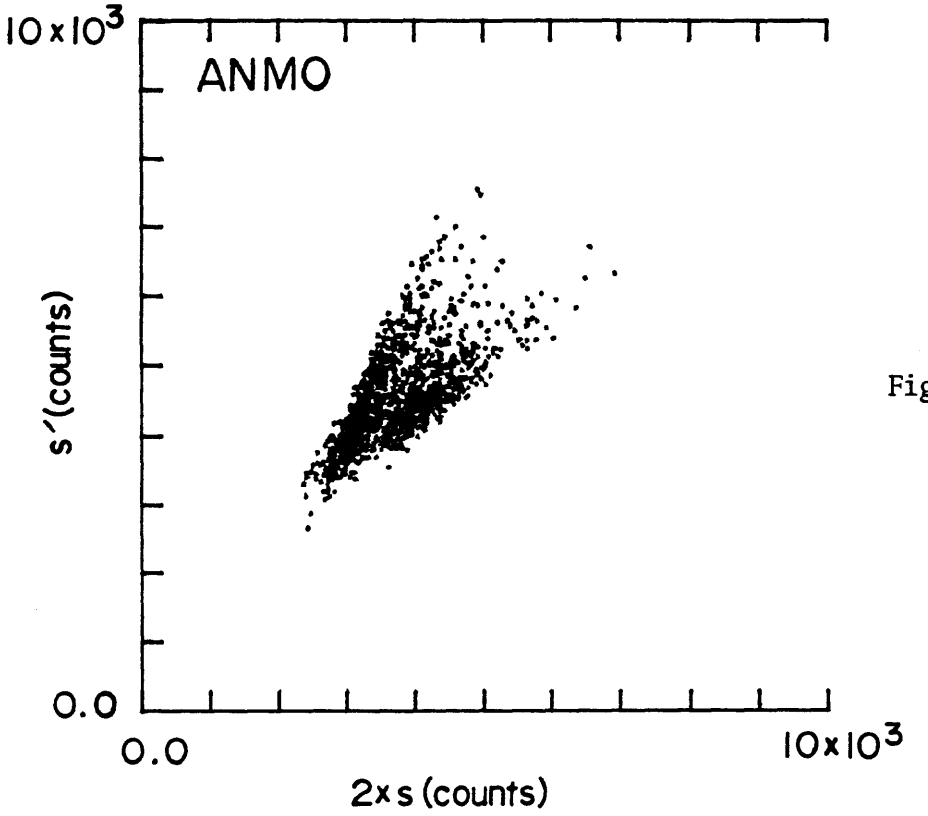

Figure 6A

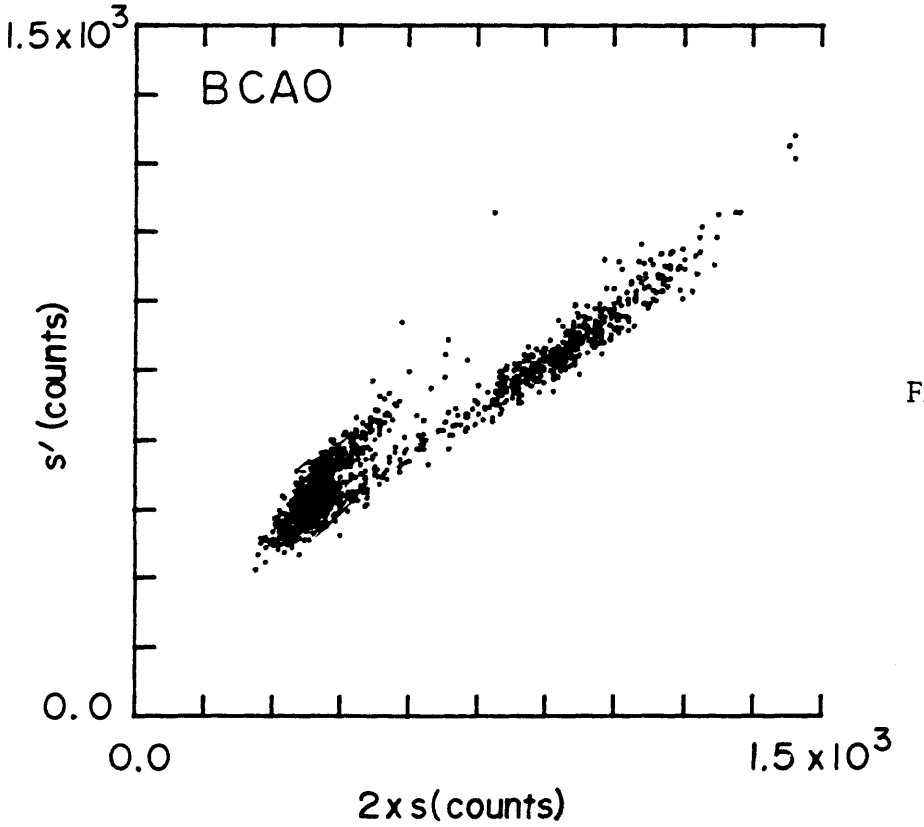

Figure 6B

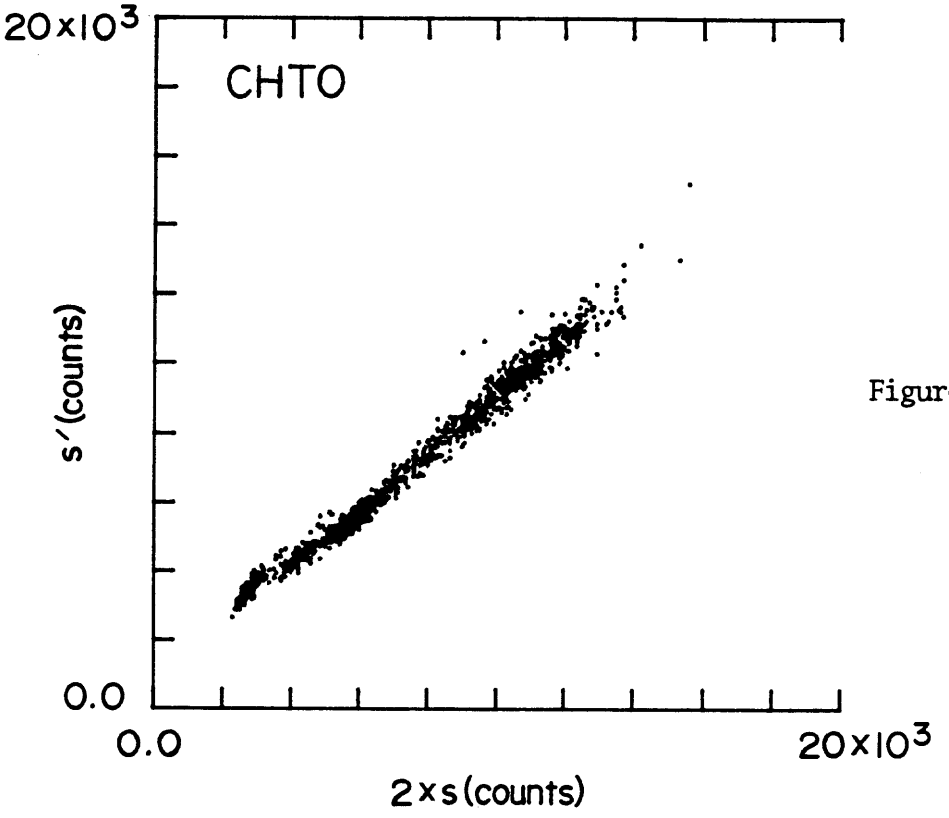

Figure 6C

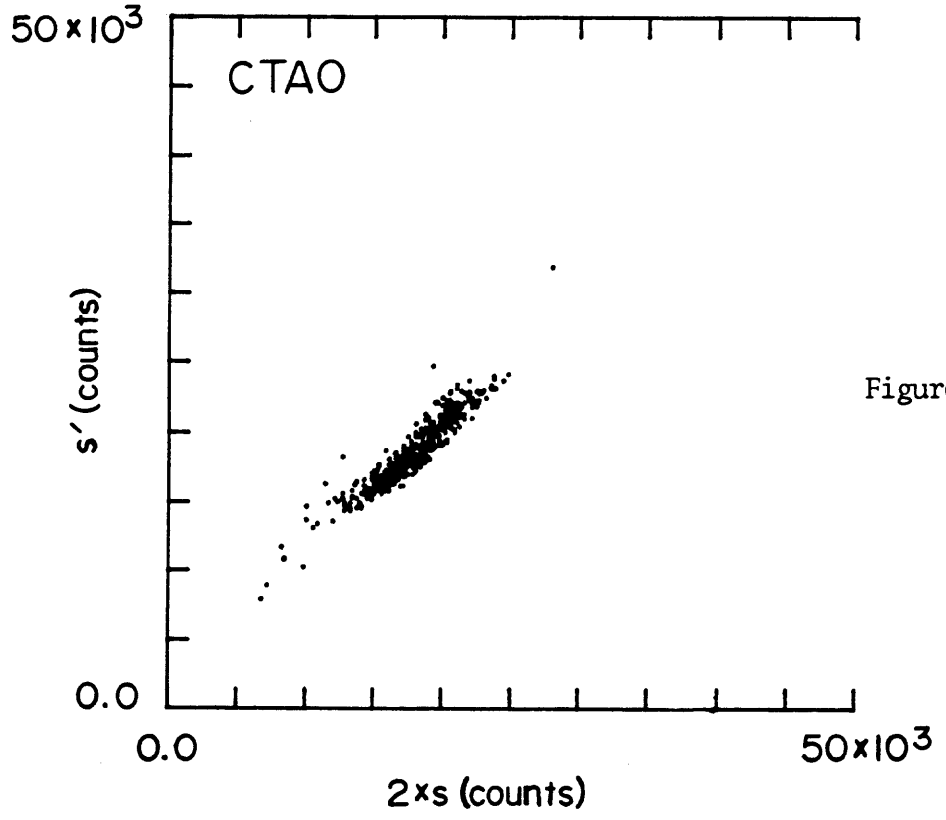

Figure 6D

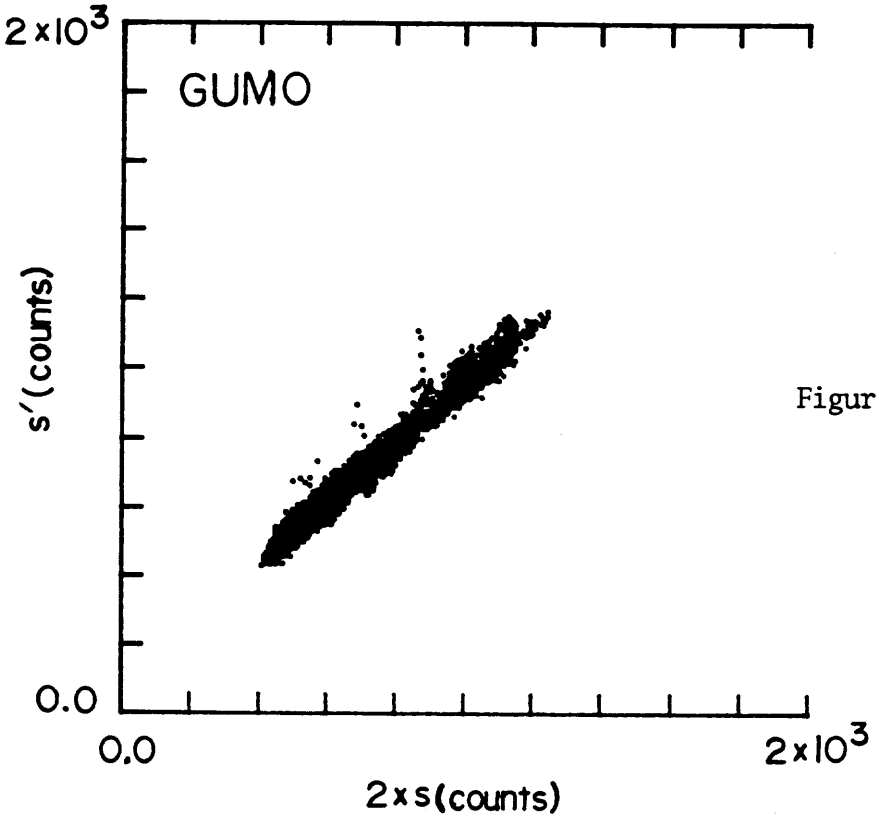

Figure 6E

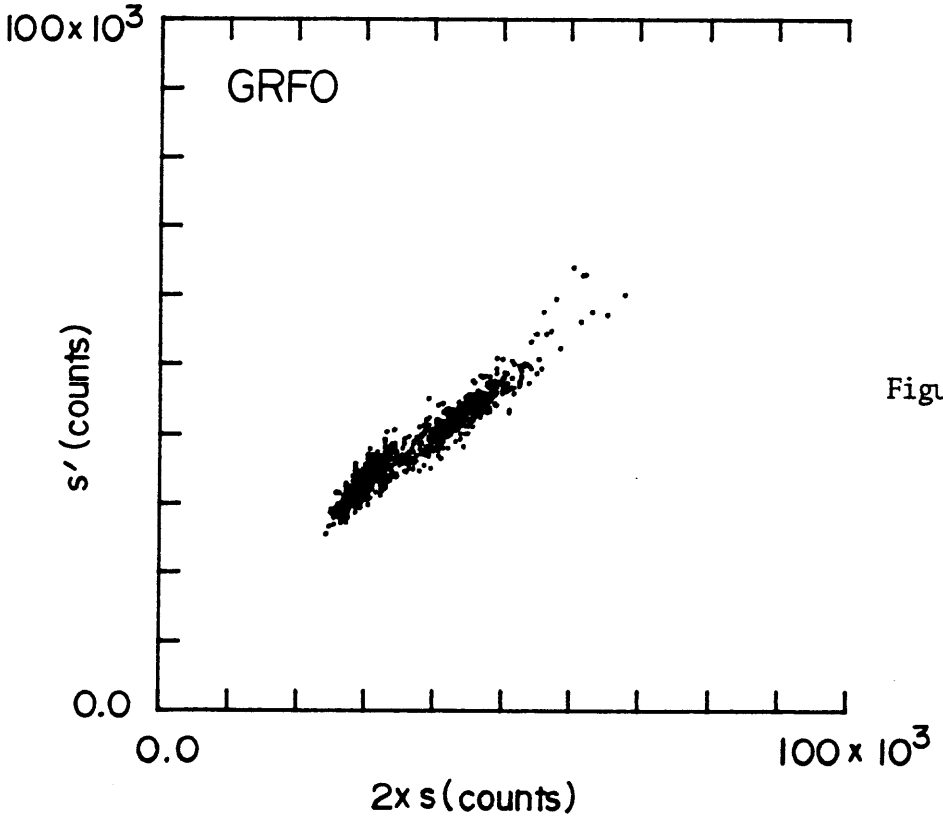

Figure 6F

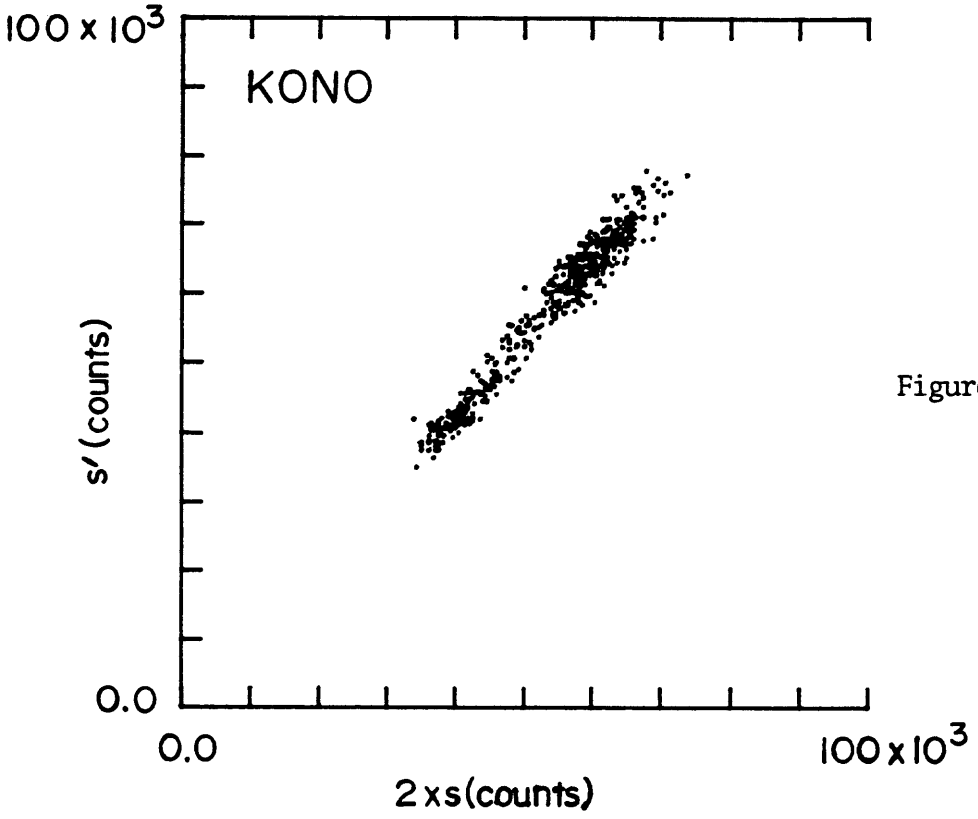

Figure 6G

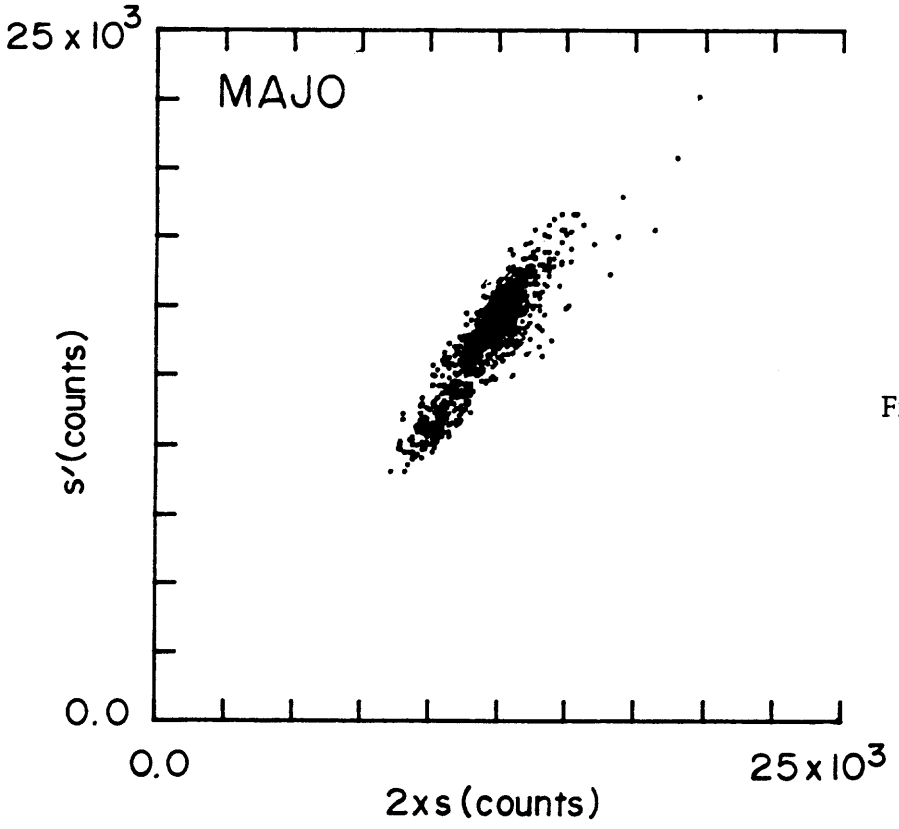

Figure 6H

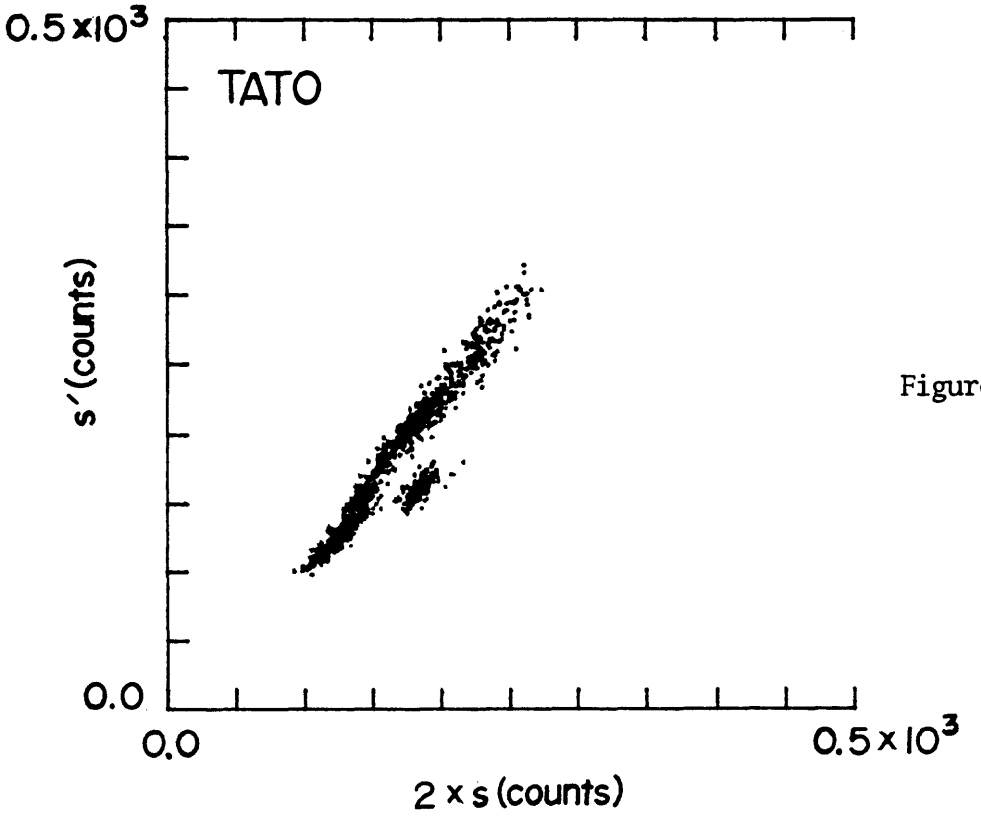

Figure 6I

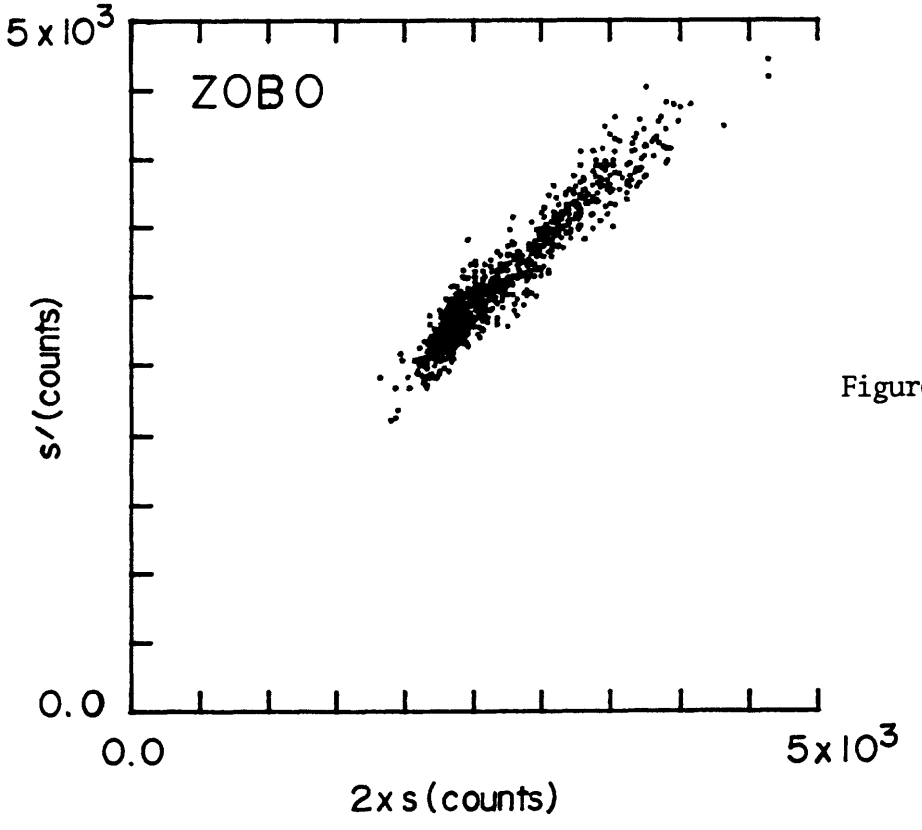

Figure 6J

Detecting an Event

Events are detected in subroutine EVENT, as described in V. above. This subroutine calls two others. One is WBUFF; it writes values into buffers when an event is judged possible. The other is IBINGO; it sets flags when an event has been declared.

For three rectified P-T values greater than the signal threshold Th2, with one of them greater than the signal threshold Th1, the detector may be thought of as searching for "Z" (or inverted "Z") shaped wave forms. For four rectified values greater than Th2, the detector may be thought of as searching for "M" or "W" shaped wave forms. Alternatively, for j rectified P-T values ($j \geq 4$) greater than Th2, the detector may be thought of as searching for changes in stationarity of the input process (complicated waveforms are easily detected). Figure 7 is a schematic of two signals that fulfill detection requirements.

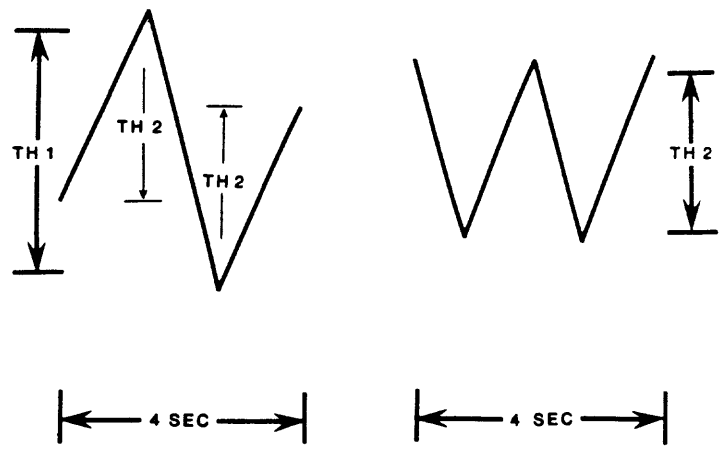

Figure 7. Illustration of two signals that meet detection requirements. The "Z" waveform on the left shows three P-T values greater than Th2, with one of them greater than Th1. The "W" waveform on the right shows four P-T values greater than Th2, with all of them less than Th1. Although these models show consecutive P-T values greater than a threshold, it is important to note that the algorithm does not require that the P-T values be consecutive.

Estimating Onset and Other Parameters of Signal

When a detection has been made, subroutine ONSETQ is called. The time (t_4, Figure 8) a P-T value first exceeded Th2 is available to ONSETQ. To estimate the onset of the signal, threshold Th3 (Th3<Th2) is used. The search (with Th3) for the onset begins two P-T values before t_4 (i.e. t_2, Figure 8) provided t_2 occurs within an appropriate time frame. (The size of the time frame is the signal period or 1.0 sec, whichever is greater.) Otherwise, the search begins at t_3, provided it occurs within the time frame. In the time frame, when the P-T value associated with t_i (i = 2, 3, 4) exceeds Th3, the onset is defined as t_{i-1}, provided $t_i - t_{i-1} < 0.5$ sec; otherwise, the onset time is $t_i - 0.5$.

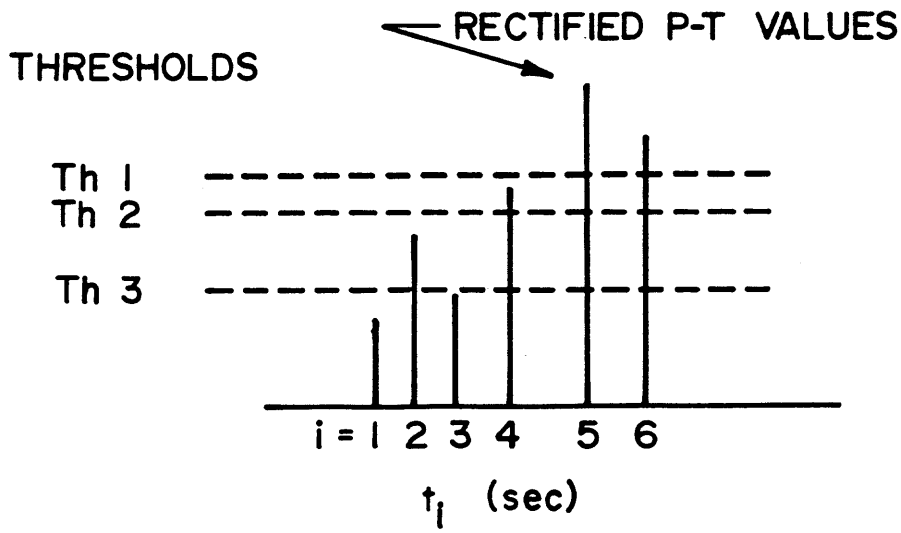

Figure 8. Timing the onsets of a signal. Six rectified P-T values are shown. The first P-T value of the detection happened at time t_4. Upon detection of the event, the algorithm looked back two P-T values before t_4 and compared them to Th3. In the illustration, the P-T value at t_2 exceeded Th3. The onset of the signal is estimated as the time of the preceding peak or trough (t_1), subject to certain restrictions imposed by time windows (see text).

Routine ONSETQ calls ONSET, a routine that estimates the quality of the onset, its polarity, and both maximum amplitude and average period of the signal (first eight P-T values). (Only a preliminary estimate of signal period was made in ONSETQ.) Traditionally, an analyst has expressed the quality of his pick purely qualitatively as e or i, depending upon whether the signal onset is judged emergent or impulsive. We have provided a more quantitative estimate of the quality of signal onset. The estimate relates to the signal-to-noise ratio (SNR) of the amplitude of the onset: the SNR is the rectified amplitude (P-T value C, Figure 9) divided by s', rounded to the nearest integer, and truncated at nine (if necessary). The SNR of each of the two rectified P-T values before the onset (A and B, Figure 9) and the SNR of each of the two rectified P-T values after the onset (D and E, Figure 9) are calculated in the same manner. The quality of the onset is expressed as these five integers (see Figure 9).

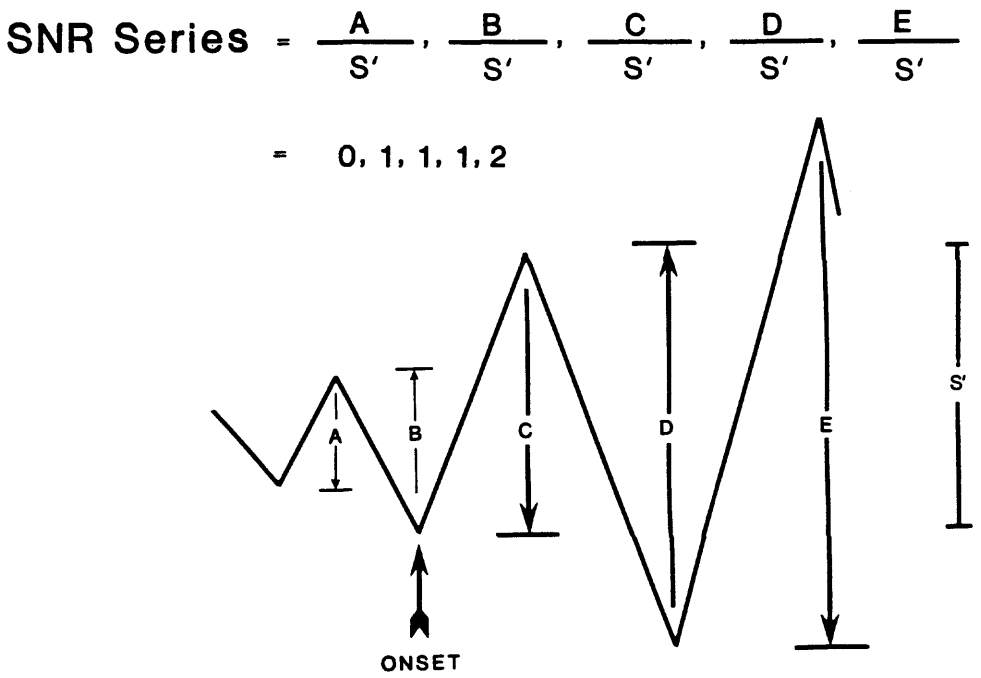

Figure 9. Illustration of what the SNR series means. In this illustration, P-T value C exceeds Th3, D exceeds Th2, and E exceeds Th1. Value A/s' is <0.5 and hence is rounded to 0.

The best possible quality would be expressed as 00999 and a poor one might be expressed as 11111, where the SNR of the estimated onset is the third integer in the series of five. The quality numbers are useful (1) to evaluate statistically the quality of the estimated onset by the relationship of each digit to s' (s is related to s' by Figure 6) and (2) as a quantification of the amplitudes of the onset of the signal (the last three quality numbers quantify the first three P-T values of the detected signal).

ONSET outputs parameters of the detection, as shown in the example below.

C 1 00125 2 360 00 36 40.62 0.432E+05 .52 0.8235E+04

where

C is the polarity of the first break (C or D)

1 is the number of P-T values looked back, relative to t_4 (0, 1, or 2)

00125 is the SNR series (the quality numbers)

2 360 00 36 40.62 is the time of onset (yr, day, hr, min, sec)

0.432E+05 is the maximum amplitude of the first 8 P-T values (counts)

0.52 is the average period of the first 8 P-T values (sec)

0.8235E+04 is s' (counts)

Using the Detector to Activate Recording of a Detected Event

Typically, the event detector is used in an operational recording system (like the SRO) to initiate the recording of short-period (SP) data on tape and to output detection messages by a printer. When this is done, it is necessary to disable the detector for a short time upon the detection of an event to prevent excessive additional detections immediately after the first one. It is also necessary to decrease the sensitivity of the detector (increase the thresholds) to prevent immediate re-triggers in the coda of events, where the P-T values are likely to be somewhat larger than those before the event. However, it is not desirable to disable new detections for too long a time or to raise detection thresholds too high, for fear of missing the onsets of significant new events.

As implemented in the SRO system, the algorithm instructs the SP tape write subroutine to write at least 196 seconds of SP data on tape, including a minimum of 20 to a maximum of 49 seconds of pre-onset data (depending on the amount of data in the SP buffer, which is partially flushed out periodically). The thresholds Th1, Th2, and Th3 are doubled (in ONSET) and event detection capability is disabled until a minimum of 49 to a maximum of 78 seconds of data after the onset have been recorded on tape. The detector is then re-enabled, so that new detections using the raised thresholds are possible during the recording of the last 98 seconds of data. Each additional detection during the recording of data on tape will cause the thresholds to again be doubled, and 196 seconds of data after the latest detection to be recorded. As the final 98 seconds of data are being recorded, PTWO calculates new estimates of s'. Hence, by the time recording on tape ceases, distinct new signals larger than the post-event background can be detected by using the post-event estimate of s'.

EXPERIENCE IN TUNING THE DETECTOR

Experience with our Model Data

The processing of models of seismic data is useful to describe the performance of a detector under idealized conditions. In particular, the models are useful for estimating false alarm rates when event-free time series are processed. To construct the event-free time series, we (1) selected four hours of short-period data (station GRFO) that did not display any events and (2) summed this series with itself 10 times, each time shifting the original by i×49 sec (i = 1, 2, 3...10). Thus, if any unseen signals were in the original series, they were significantly attenuated in the sum. The auto-correlation for the sum series (Figure 10) demonstrates the absence of strong periodic properties, and the probability density estimate for the series (Figure 11) appears approximately normal. To describe the dispersion of the filtered (F1) 4-hour P-T series, each P-T value of the filtered series has been rectified and normalized by dividing it by the associated value of the estimate of s (s'/2). The cumulative distribution of these P-T values is shown in Figure 12. There are an average of 3.6 P-T values/sec.

The model time series (sum series, above) has been used to demonstrate the false alarm rate for different settings of the parameters. To summarize, there are five sets of major parameters that may be varied to tune the detector:

1. Threshold factors Xth1, Xth2, and Xth3

2. Window length.

3. Filhi and Fillo, these are lengths of subwindows of the window length.

4. Number of P-T values, greater than the thresholds, that are required for detection. There are two subalgorithms here: Subalgorithm 1 declares a detection if one P-T value is greater than Th1 and any two other P-T values are greater than Th2; subalgorithm 2 declares a detection if any j ($j \geq m$) rectified P-T values are greater than Th2 but are less than Th1. Generally $m \geq 4$.

5. Filter type. F1 (Figure 2) was used for the tests.

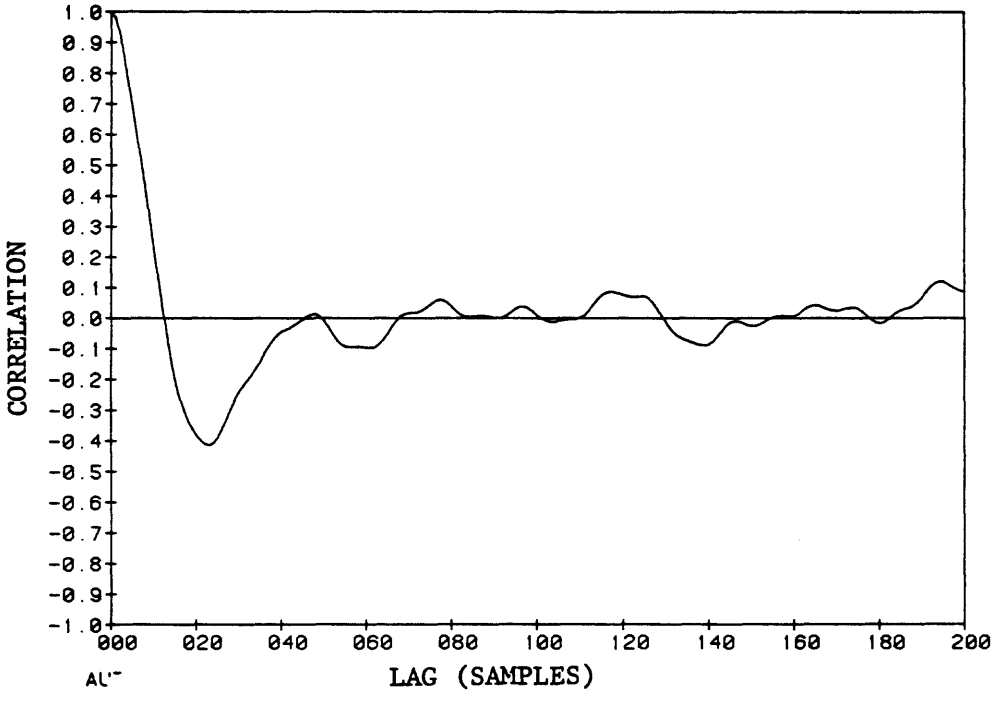

Figure 10. Auto-correlation of a sample of the time series that was used to estimate false alarm rates for different settings of the detector. The sample is 245 sec long, 20 sps.

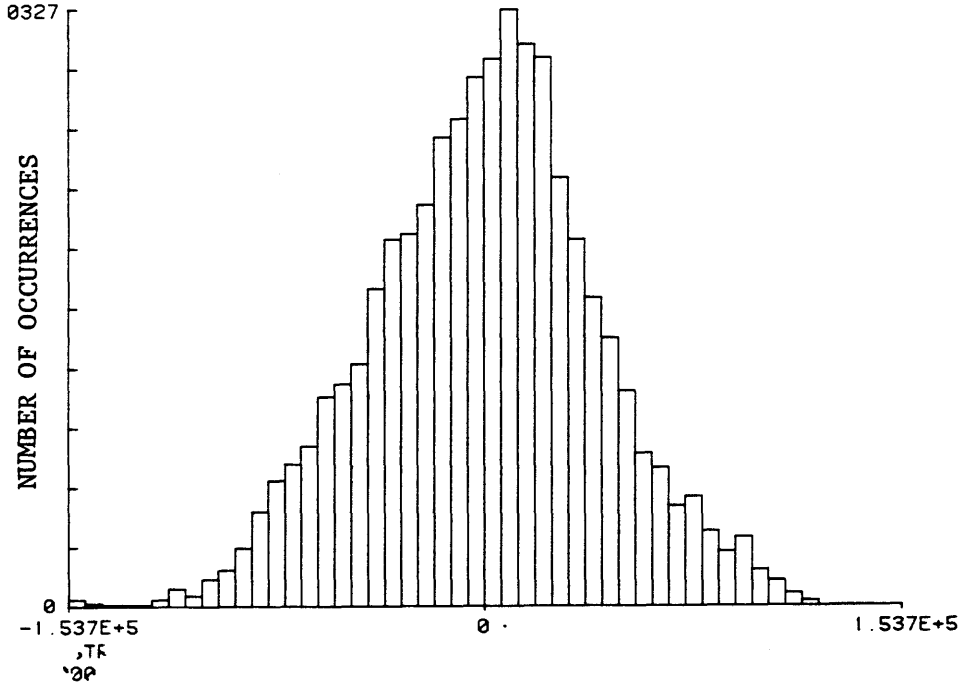

Figure 11. Histogram of values of the digitized data that were used to estimate the auto-correlation function.

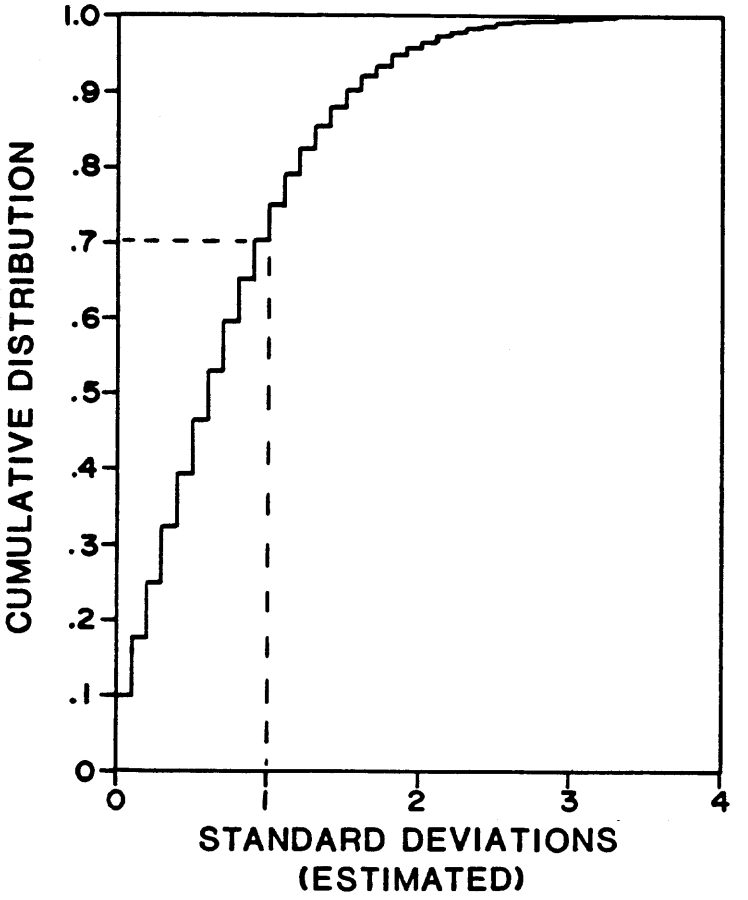

Figure 12. Cumulative distribution of the normalized P-T values as a function of the estimate of s (s'/2).

For Figure 13, parameters varied are Xth1, Xth2, and m. The window length (4 sec), Filhi (0.2 sec), Fillo (2 sec), and Xth3 (1.0) are held constant. For two values of m (4 and 6), Figure 13 shows false alarms as a function of Xth1 and Xth2, and for each value of the variables, two false alarm rates are shown. One (the upper one) is for subalgorithm 1 and the other (the lower one) is for subalgorithm 2. For instance, for m = 4, Xth1 = 2, Xth2 = 1.0 (Figure 13), there are 1.0 FA/hr due to subalgorithm 1 and 9.2 FA/hr due to subalgorithm 2. For m = 6, the same settings of the thresholds produce only 0.8 FA/hr[1], all due to subalgorithm 1. Thus, increasing the number of cycles required for the detection (m = 4 corresponds to two cycles and m = 6 corresponds to three cycles, steady state) significantly reduces the false alarm rate. The illustration demonstrates that the most pronounced difference in the rates (i.e., those for m = 4 vs those for m = 6) is for the large values of Xth1 paired with the small values of Xth2, as one might expect.

Although they can be employed to produce large changes in the false alarm rate, the subwindows Filhi and Fillo normally fine tune the detector. The value of the subwindows is important (in reducing false alarms) mainly when small values of Xth2 are employed. This is demonstrated in Figure 14: For Xth2 = 1.0, changing Filhi from 0.1 sec to 0.5 sec does not appreciably reduce the false alarm rate; however, for Xth2 = 0.8, changing Filhi as above reduces the false alarms by about 20% for m = 4 and substantially for m = 6. Changing the value of subwindow Fillo affects the false alarm rate somewhat more substantially, as demonstrated by Figure 15 where the false alarm rates are displayed for Fillo = 1.0 sec to 4 sec in 0.5 sec increments. For both values of m and for the larger value of Xth2 (1.3), reducing Fillo from 4.0 sec to 1.0 sec decreases the false alarm rates by about 50%, and for the smaller values of Xth2, there is a much larger reduction. However, there is a price associated with the reduction in the false alarms. For Filhi, increasing it to too large a value likely will cause high frequency signals of about 1-2 cycles of duration to be missed. For Fillo, decreasing it to too small a value likely will cause onsets of emergent events to be timed late and small emergent signals to be missed altogether.

The data of Figures 13, 14, and 15 have been used to demonstrate the false alarm rates when artificial noise is input to the algorithm. How well the values of Figures 13, 14, and 15 estimate false alarm rates for other filtered (F1) series very likely will depend largely on the correspondence between the distributions of their normalized P-T series to that of the artificial noise and on the correspondences of the average number of P-T values per unit time.

[1]Differences between experimental measurements of 0.2-0.3 FA/hr translate to 1 FA in the 4-hr sample. One likely source of this small variation is the start-up of processing. Another likely source is a small increase in s' caused by including values in its estimate (for m = 6) that were declared signals for m = 4.

Xth2	m = 4						m = 6						Xth2
1.5	0.8 0.0	0.8 0.0	0.8 0.0	0.5 0.0	0.2 0.0	0.0 0.2	0.8 0.0	0.8 0.0	0.8 0.0	0.5 0.0	0.2 0.0	0.0 0.0	1.5
1.4	0.8 0.0	0.8 0.0	0.8 0.0	0.5 0.0	0.2 0.0	0.0 0.2	0.8 0.0	0.8 0.0	0.8 0.0	0.5 0.0	0.2 0.0	0.0 0.0	1.4
1.3	1.8 0.0	1.8 0.0	1.8 0.0	1.0 0.0	0.5 0.0	0.0 0.2	1.8 0.0	1.8 0.0	1.8 0.0	1.0 0.0	0.5 0.0	0.0 0.0	1.3
1.2	4.8 0.2	4.0 0.5	2.8 0.8	1.8 1.5	0.5 1.8	0.2 1.8	4.8 0.0	4.0 0.0	2.8 0.0	1.8 0.0	0.5 0.0	0.0 0.0	1.2
1.1	8.8 0.2	6.5 0.8	4.0 1.5	2.8 2.2	1.0 2.5	0.8 2.5	8.8 0.0	6.5 0.0	4.0 0.0	2.8 0.0	1.0 0.0	0.5 0.0	1.1
1.0	12.2 5.2	8.5 6.5	5.8 8.0	3.5 8.5	1.2 9.2	1.0 9.2	12.2 0.0	8.5 0.0	5.5 0.0	3.5 0.0	1.2 0.0	0.8 0.0	1.0
	1.5	1.6	1.7	1.8	1.9	2.0	1.5	1.6	1.7	1.8	1.9	2.0	
			Xth1						Xth1				

Figure 13. False alarm rates (FA/hr) for different settings of threshold factors, Xth1, Xth2, m = 4 and m = 6. Window length = 4 sec, Filhi = 0.2 sec, Fillo = 2.0 sec, Filter is F1. For each setting of Xth1 and Xth2, two experimental values are shown. The upper one is the false alarm rate for subalgorithm 1, and the lower one is the false alarm rate for subalgorithm 2.

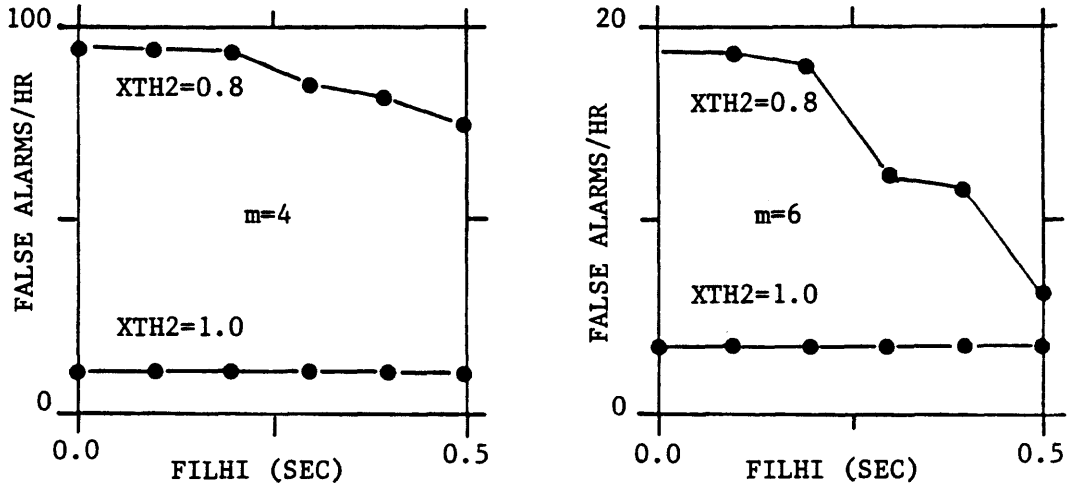

Figure 14. The effect of changing Filhi. Two different settings of Xth2 are shown. The filter is F1, Win = 4.0 sec, Fillo = 2.0 sec, Xth1 = 1.8. The illustration on the left shows results for m = 4 (4 P-T values greater than Th2), and the one on the right shows results for m = 6.

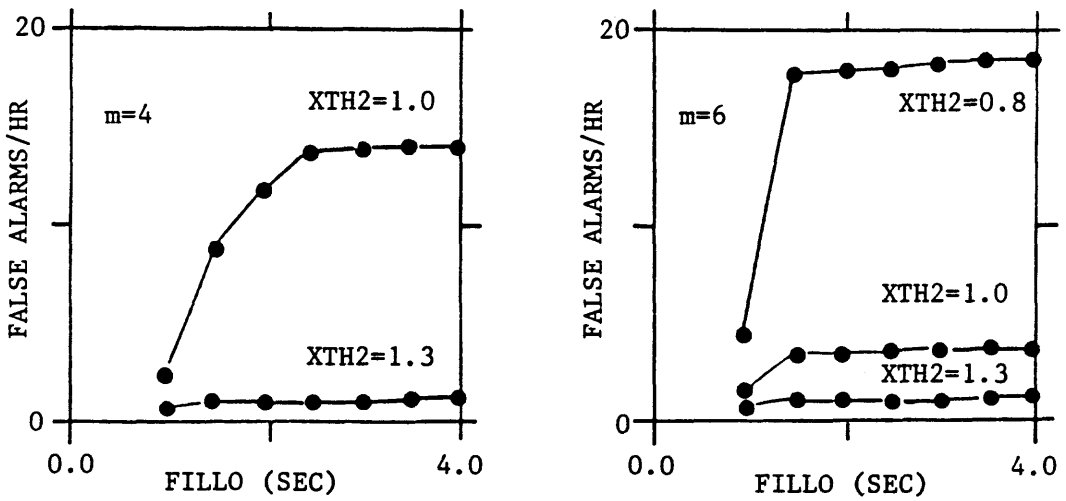

Figure 15. The effect of changing Fillo. Various different settings of Xth2 are shown. The filter is F1, Win = 4.0 sec, Filhi = 0.2 sec, Xth1 = 1.8. The illustration on the left shows results for m = 4 (4 P-T values greater than Th2), and the one on the right shows results for m = 6.

Experience with Blandford's Model Data

All references in this section are of the report, Blandford et al., 1981. Blandford and his coworkers generated two sets of model data. One was constructed from real seismic data recorded at Pinedale, Wyoming and the other was constructed from the same type of data recorded at NORSAR. The results reported here are for the Pinedale data (20 sps), we have not yet processed the NORSAR data (10 sps). The method of constructing the artificial noise was described in their report (1981, p. II-2).

> A 4096 point data window is read in and checked for signals and data dropouts. If there are none, then the data are 5% cosine tapered, Fourier transformed, the phase at each frequency is assigned randomly, and the data are transformed back to the time domain. Then the noise is tapered with a full 50% taper and added to the previous window (which has also been tapered) with a 50% offset. Because the sum of the two 50% tapers does not yield a constant root-mean-square amplitude this fact is corrected for by multiplying the summed time series by an analytical function of time. The successive overlapped portions are read out to the final noise tape leading to a continuous random noise field whose spectrum and amplitude vary smoothly as does the true noise field.
>
> The next step in the process is to add in the signals...

Thirty-one signals recorded at Pinedale were added to the noise in the following manner by them (1981, p. III-2).

> The signal windows are two minutes long and the P starts were centered in the window. Thus, there is one minute of noise in front. The signal windows were tapered with a 25% cosine taper to avoid abrupt starts and stops, then added to the noise. Each signal was added to the noise four times, first with the maximum equal to 1/2 the raw amplitude maximum in the 10 minute window, and then, at 1/4, 1/8 and 1/16 the amplitude. Thus, with the 31 signals...there are a total of 4×31 = 124 10-minute (12000 points) windows. In each window a signal begins at the 9th minute (point 10800). The data are continuous from window to window.

These data were processed with a number of different algorithms by them and others (1981). A detection was scored if a trigger occurred within 30 sec after the signal had been added (1981, p. II-3).

To enable analyzing the different algorithms on an equal basis, they (1981) devised a scheme to produce comparability by using a reference level. First, they (1981, p. 12) calculated the false alarm rate (FA/hr) and the "corrected number of detections."

> A corrected number of detections is...determined by subtracting an estimate of the number of detections which are probably false alarms. This is determined approximately as the false alarm rate times the total signal windows in which a detection is not expected to occur. These signal windows are taken to be those in the 3rd and 4th S/N level, plus 1/2 of those in the 2nd S/N level.

There are a total of 124 windows and each signal window is (0.5/60) hours; and, as explained above, they assume that there are no detectable signals in (2.5/4) of the windows. Thus the number of detections that are probably false alarms is (2.5/4)×124×(0.5/60)×FA/hr (1981, p. 12). To produce comparability they subtract 46 events from the "corrected number of events", where 46 is the number of events they would expect to detect (all of those of the 1/2 amplitude level plus one-half those of the 1/4 amplitude level). In comparing two different detectors at a given false alarm rate, if detector A detected all of the events 6 dB lower than the other (B), Blandford and his coworkers would consider detector A 0.3 m_b more sensitive than B. Because there are ∿30 signals in each 6 dB (0.3 m_b) level, if detector C detected one more event than A, they would consider C 0.01 m_b (0.3/30 = .01) more sensitive than A (see 1981, p. 13). They plot the log of the false alarm rate as a function of this relative magnitude (see Figure 16 where an abstract of their Figure 1 is displayed).

Table 1 shows the results of processing the Blandford data with our routine ("ASL on line") detector. (It is essentially the detector that was implemented with filter F1 on the SRO and ASRO stations, May 3, 1983.) The table displays the effect of changing values for the parameters Filhi, Fillo, Xth1, and Xth2. Whereas small changes in Xth1 or Xth2 can significantly affect the operation of the detector (as was demonstrated in the previous section), small changes in Filhi or Fillo affect the operation only moderately.

For Filhi = 0.2 sec and Fillo = 2.0 sec (typical operational settings) Figure 17 shows the performance of our detector compared with that of the Seismic Data Analysis Center (SDAC), Alexandria, Virginia; the SDAC detector is implemented with a 6-pole 1-2 Hz band-pass filter. The performance of these two detectors is very similar. This correspondence demonstrates that our detector, implemented with a simple sum and difference filter (Figure 2), can compete successfully with a detector that uses a sharp band-pass filter. Also, the performance of two analysts (Blandford et al., 1981) is shown in Figure 17. The performance of the detectors seems approximately the same as that of the two humans. Thus, for the false alarm rates tested, the detectors appear to be competing successfully with the humans.

Run 7 (Table 1) is for operational settings of the detector; the run gave one false alarm (∿1FA/day). To demonstrate again the performance of the detector, Figure 18A shows representative examples of events detected on Run 7 and Figure 18B shows representative examples of the largest events not detected on this run. On each illustration, four sets of two traces are shown; the upper trace of each set is the data supplied by Blandford, and the lower trace is the same data filtered by F1. For the events detected, the output onset time and direction of first break are indicated. Normally, if we are able to pick the onset of detected events with confidence, approximately the same onset is picked by the algorithm.

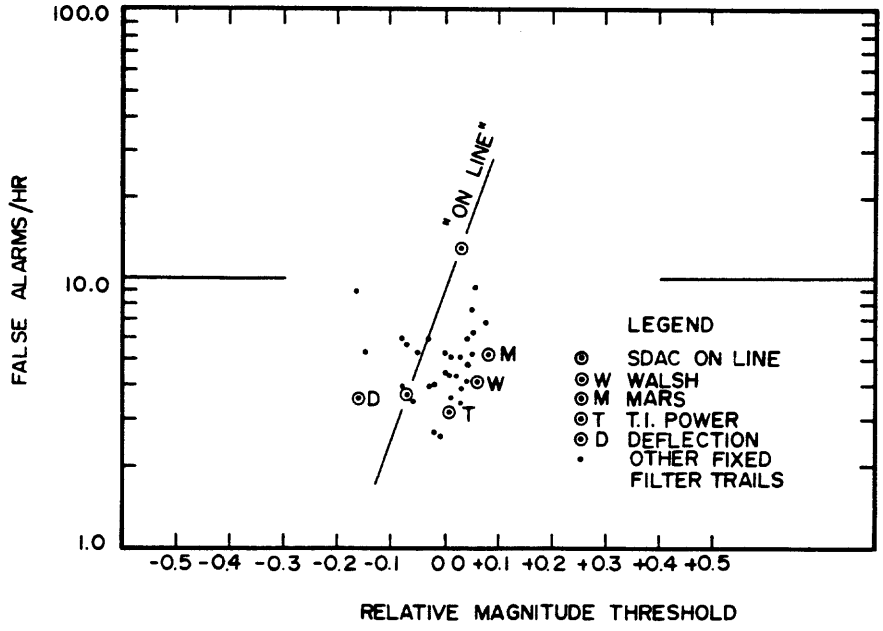

Figure 16. A generalization of Figure 1 of Blandford et al. (1981), showing the results of experiments that used a single filter for all of the signals, as opposed to a different filter for each signal. At a given false alarm rate, better detector performance is for the larger relative magnitude values.

RUN	Filhi Sec	Fillo Sec	Xth1	Xth2	Win Sec	Total Events	False Alarms	False Alarms/hr	Corr. Events
1	0.1	2.0	1.7	1.0	4.0	56	281	14.31	47
2	0.2	2.0	1.7	1.0	4.0	56	258	13.14	48
3	0.3	2.0	1.7	1.0	4.0	52	205	10.44	45
4	0.1	2.0	1.75	1.1	4.0	40	68	3.46	38
5	0.2	2.0	1.75	1.1	4.0	39	68	3.46	37
6	0.3	2.0	1.75	1.1	4.0	37	53	2.70	35
7	0.2	2.0	2.0	1.5	4.0	10	1	0.05	10
8	0.2	2.0	1.75	1.25	4.0	23	13	0.66	23
9	0.2	2.0	1.75	1.375	4.0	19	5	0.25	19
10	0.1	3.0	1.75	1.1	4.0	45	98	4.99	42
11	0.2	3.0	1.75	1.1	4.0	44	98	4.99	41
12	0.3	3.0	1.75	1.1	4.0	38	80	4.07	35

Table 1. Results of processing the Blandford model data with our on-line detector (detection declared if one P-T value greater than Th1 and two other P-T values greater than Th2, or if four P-T values greater than Th2 only).

The signal periods for the Pinedale data were reported (1981) 1.0 sec or longer. Because filter F1 peaks at 0.5 sec, the performance described here (Figure 17) might realistically be considered a conservative estimate of the capabilities of the detector.

The Blandford data have also been processed by using m = 6 with the other parameters the same as Table 1. The results are shown in Table 2. Although setting m = 6 materially reduces the number of false alarms for the smaller values of Xth2, for the larger values (runs 7 and 9) there was not any measurable reduction. The large reduction in false alarms was accompanied by a significant decrease in the number of events detected (relative to those of Table 1), so that the net effect of setting m = 6 was a reduction in the performance. The reduction in performance is exemplified by runs 7 and 9: although the false alarms for m = 4 and m = 6 are the same, m = 6 produces 10-20% fewer detections. Very likely, the capability of setting m = 6 can be better exploited when the signal periods are smaller than those of the Pinedale data. One might search for local or regional events with m>4.

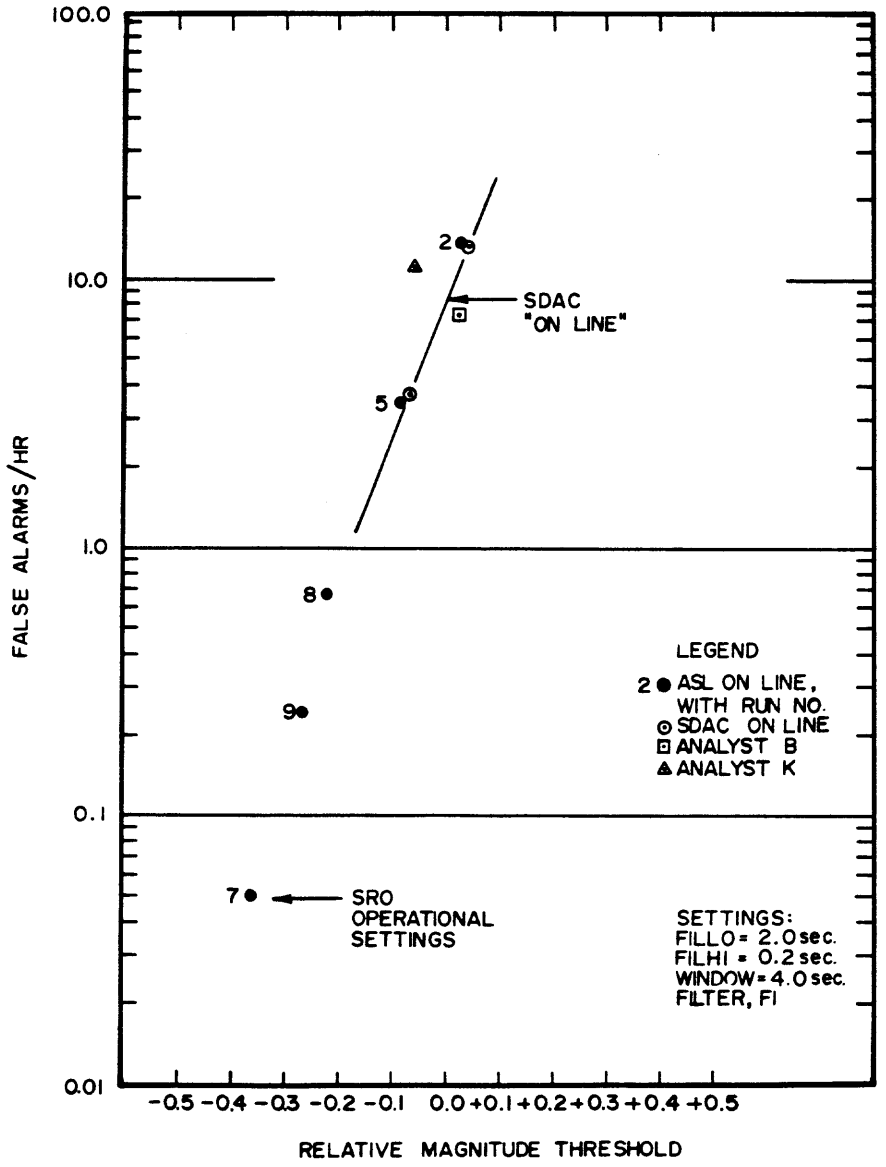

Figure 17. Results of our processing Blandford's (1981) Pinedale tape compared to results of the SDAC on-line detector and results of two analysts. Run numbers correspond to those of Table 1.

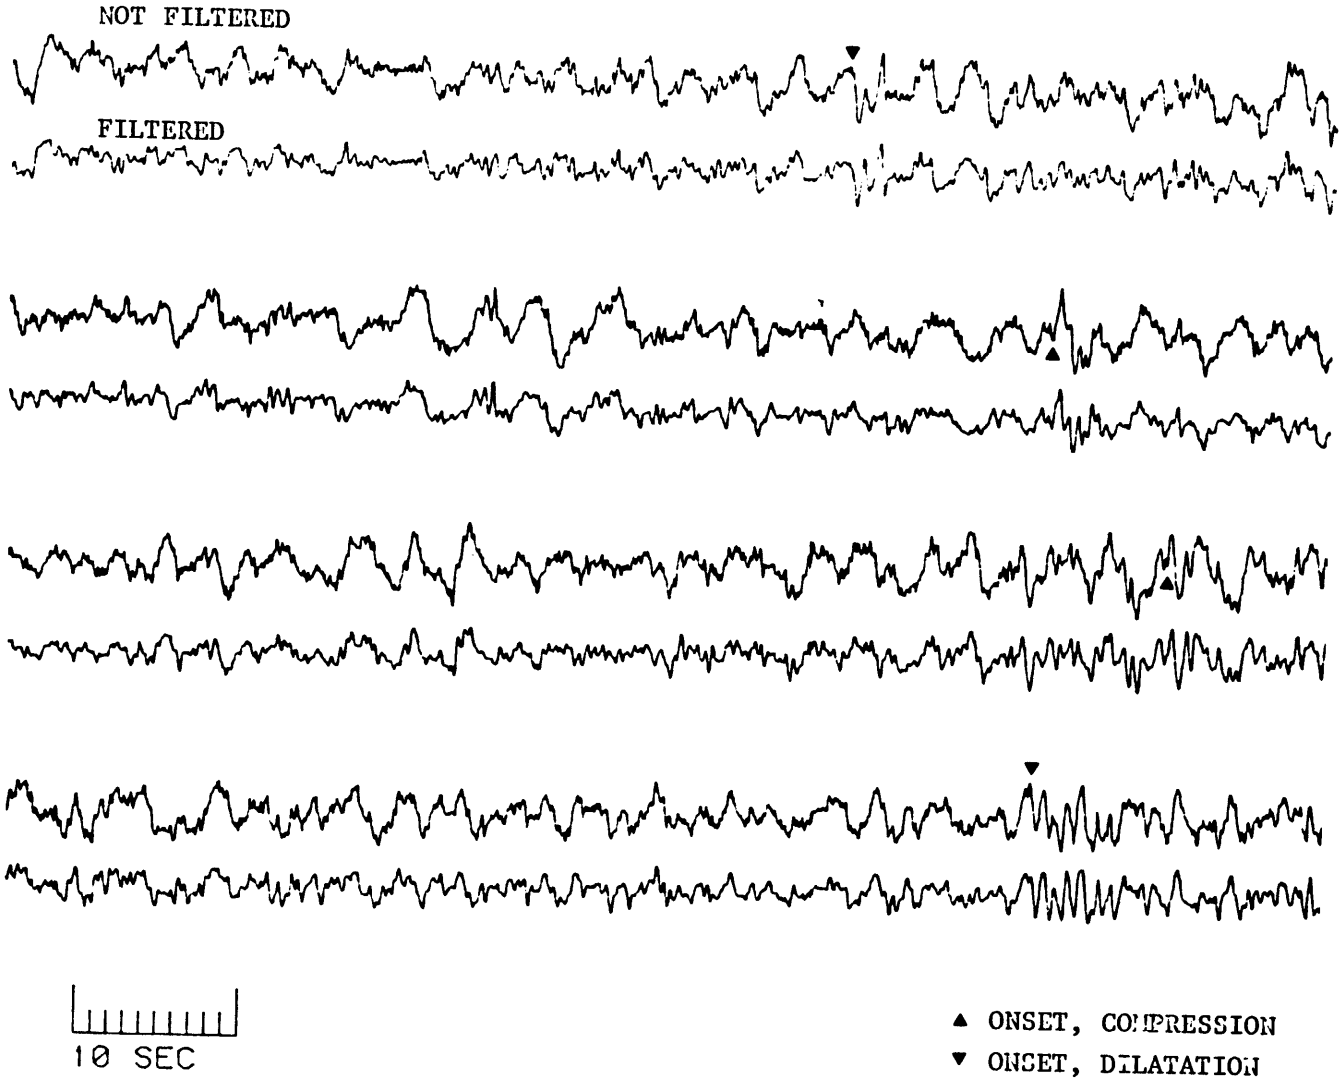

Figure 18A. Examples of signals detected on Run 7 (Table 1). The detector settings for Run 7 produce ∿1FA/day. The output onset times and direction of first motions are indicated.

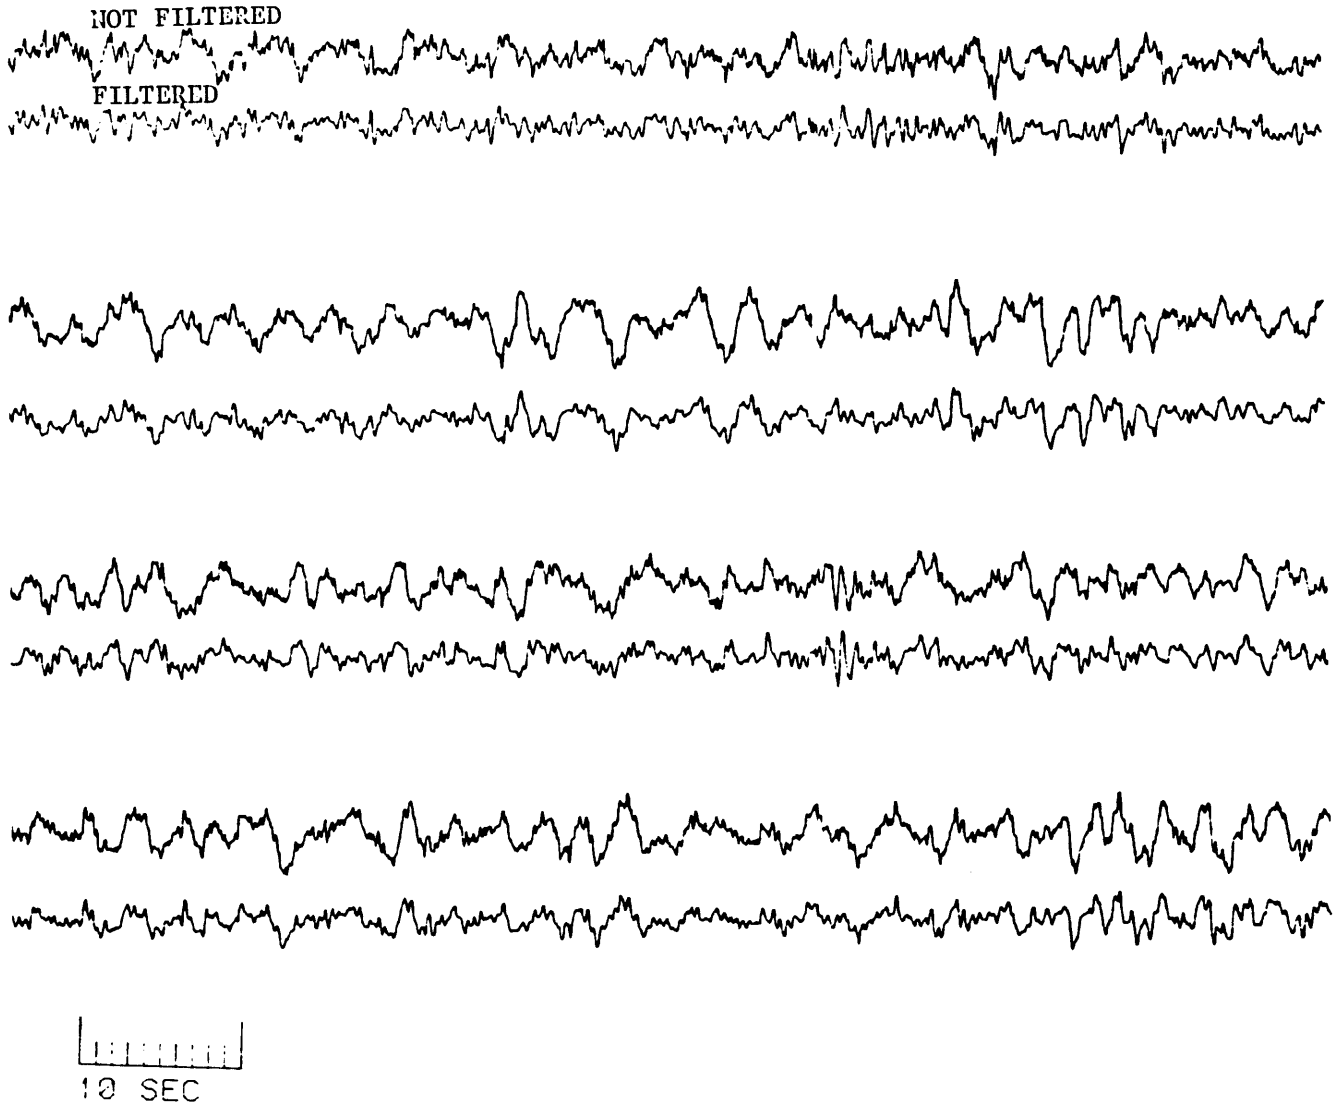

Figure 18B. Representative examples of the largest signals that were not detected on Run 7 (Table 1). The detector settings for Run 7 produce ∿1FA/day. The signals occur in the right portion of the illustration, noise is in the left portion.

RUN	Filhi Sec	Fillo Sec	Xth1	Xth2	Win Sec	Total Events	False Alarms	False Alarms/hr	Corr. Events
1	0.1	2.0	1.7	1.0	4.0	39	67	3.41	37
2	0.2	2.0	1.7	1.0	4.0	37	66	3.36	35
3	0.3	2.0	1.7	1.0	4.0	34	56	2.85	32
4	0.1	2.0	1.75	1.1	4.0	26	26	1.32	25
5	0.2	2.0	1.75	1.1	4.0	28	26	1.32	27
6	0.3	2.0	1.75	1.1	4.0	24	22	1.12	23
7	0.2	2.0	2.0	1.5	4.0	8	1	0.05	8
8	0.2	2.0	1.75	1.25	4.0	20	11	0.56	20
9	0.2	2.0	1.75	1.375	4.0	17	5	0.25	17
10	0.1	3.0	1.75	1.1	4.0	33	43	2.19	32
11	0.2	3.0	1.75	1.1	4.0	33	41	2.09	32
12	0.3	3.0	1.75	1.1	4.0	30	38	1.94	29

Table 2. Results of processing the Blandford model data with m = 6 (detection declared if one P-T value greater than Th1 and two other P-T values greater than Th2, or if six P-T values greater than Th2 only).

Experience with the October 1980 Data

During October 1980, short-period data were recorded full-time for two weeks at twelve of the operating SRO and ASRO stations. These data have been used to obtain operational settings for each of the stations. Requirements of the network-day tape operations at ASL impact on the detector settings: The day tape operations cannot easily accommodate more than an average of 30 detections per station per day. Although it is relatively easy to tune the detector to produce essentially zero microseism-related false alarms per day, yet pick all (or almost all) of the clearly recorded events seen on the seismograms, it is much more difficult to tune it to reject prominent cultural noise, yet still pick small events satisfactorily. The requirement to reject a substantial amount of cultural noise has degraded somewhat the performance of the detector at several of the stations, particularly BCAO, BOCO, and GRFO. GRFO has been affected more than the others, probably because many of the signals there are very emergent.

The best operational performance of the detector, relative to an analyst, is exemplified by CHTO. Intermediate operational performance is exemplified by BCAO. Substandard performance is seen at GRFO. So that the reader can judge what we mean by "best," "intermediate," and "substandard," a representative seismogram, together with the corresponding print-out from the detector, is presented for each of the classes (Plates 1, 2, 3)[2]. (It is important to note that we are processing all of the data with our general purpose filter, F1). Even under the severe conditions imposed (outlined above), the detector usually picks all of the clear events displayed on a normal[3] SRO or ASRO short-period seismogram. Quantitatively, we would expect event A (Figure 19) to be in the class of events detected with a 0.5 or greater probability, and event B (Figure 19) in the class detected with a 0.9 probability, using standard operational settings.

The former SRO detector operated poorly at some of the stations, such as GUMO, where the amplitudes of the background noise often changed substantially. In such instances, many false alarms might occur. The new detector is able to track large changes in the amplitude of the background yet produce only a few microseism-related false alarms per day, see Plate 4.

[2]A description of the meaning of the output is given on page 17.

[3]Except for NWAO, all of the helicorders have approximately the same frequency response as the digital data. This response peaks at 0.4 sec. At NWAO, the frequency response for the helicorder peaks at 0.15 sec. At this station, small clear high-frequency events displayed on the helicorder may not be detected. On the other hand, the detector picks lower frequency signals that are not clearly displayed on the helicorder.

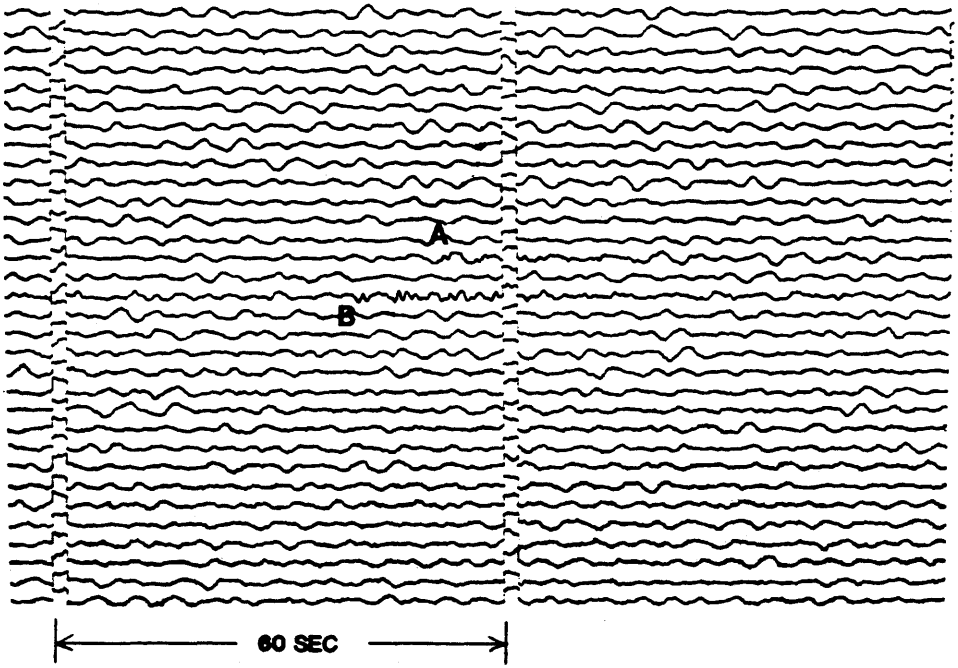

Figure 19. A type of signal judged detected with 0.5 probability (A) and one judged detected with 0.9 or greater probability (B), microseism-related false alarm rate $\sim 10^0$ per day. The station is ANMO; it is in the set of the better stations.

SPEED AND SIZE OF THE ALGORITHM

The program was developed and tested on a Digital Equipment Corporation PDP 11/34 Computer. Although the algorithm was designed for implementation by simple shifts, additions, and comparisons, on this machine simple Fortran arithmetic instructions were used. No effort was made to optimize the code for speed. After development, the algorithm was implemented on the Data General NOVA 1200 computer of the SRO. On this machine, the algorithm was programmed without using any floating-point operations. The following gives a description of the speed and size of the algorithm on these two different computers. All speeds are for operational settings: Xth1 = 2.0, Xth2 = 1.5.

1. Machine: PDP 11/34 (DEC)

   Operating System: RSX-11M

   Special Equipment: Floating Point Processor

   Language: Fortran 4 Plus, ∿500 lines of code required

   Task Image Size: 20,640 16-bit words total

   Breakdown:
   a) Fortran-callable SRO tape reading routine = 9,620 words
   b) Data buffer in compiled Fortran program to hold one SRO record (990 values) = 1,980 words
   c) Remainder of compiled Fortran program (including code, buffers, constants, and variables) = 9,040 words.

   Speed: 30 × real time, reading a single SRO SP channel, 20 samples per second (includes overhead of reading tape, error checking, and buffering the input record).

   40 × real time, reading data from disk

2. Machine: NOVA 1200 Jumbo (Data General) with magnetic core memory

   Operating System: SRO real-time operating program

   Special Equipment: None

   Language: Assembly

   Size: 4,000 16-bit words

   Breakdown: Data buffers = 700 words
   Code, constants = 3,300 words

   Speed: 25 × real time (former SRO detector was 10 × real time)

WORK IN PROGRESS

Four significant modifications to the detector are being made and tested. First, a new algorithm is being developed to pick secondary phases in the coda of events. Plate 5 demonstrates the progress we have made with this improvement. At present, we are working toward reducing the number of triggers in the coda of large events, without sacrificing detection of the clear phases in the coda of small events.

Second, several different window lengths, together with several different filters, are being tested on the Pinedale tape of Blandford. Work to date appears encouraging: by using a 1-2 Hz recursive filter, and by shortening the window length and making minor modifications to the algorithm, we can obtain results a bit superior to those of the best detectors of Figure 16. We need to test (1) the generality of these results, (2) the effect of these modifications on the fidelity of picking onsets and polarity, and (3) the effect of substituting different filters for the 1-2 Hz recursive filter.

Third, tests are being conducted on the feasibility of using the parameters of a detection (period, amplitude, quality numbers) to identify and reject (on-line) cultural noise.

Fourth, an algorithm is being tested to identify teleseisms (by using the signal period) and to make recording length for each teleseism a function of period and amplitude of the signal. The purpose of the modification is (1) to ensure recording of secondary phases (such as S and P'P') even if their spectrums are not in-band to the filter and (2) to ensure recording of small signals that might go undetected in the coda of large events.

SUMMARY

A computer program has been developed and tested to detect short-period seismic events. It operates on a single channel of digital data. All operations can be performed by simple sums, shifts, and comparisons; hence, the algorithm can be easily implemented even on machines that do not have floating point hardware.

Figure 1 shows the flow of the program. Briefly described, the input time series is filtered with a simple FIR filter (Figure 2). The output of the filter is processed to find relative maximums and minimums (the peaks and troughs) of the series. The maximums and minimums are differenced to obtain a new time series (the P-T series). The remainder of the algorithm proccesses the P-T series only. The dispersion (variability) of the P-T series is estimated by a comparison scheme: the maximum of 20 successive rectified P-T values is found, and an average of 16 such maximums is made. For normally distributed P-T values of zero mean, this average (s') is an estimate of twice the sample standard deviation of the P-T series.

Three thresholds are employed and all are a function of s'. Two of them, Th1 and Th2 (Th1>Th2), are used for detecting events, and the other (Th3, Th3<Th2) is used for timing the onsets of events. In the typical operating configuration, a detection is declared if (1) one rectified P-T value exceeds Th1 and two other rectified P-T values exceed Th2, or if (2) four rectified P-T values exceed Th2. Although complicated wave forms are easily detected, the algorithm can be thought of as searching for "Z" shaped or "W" or "M" shaped waveforms (see Figure 7). When a detection has been made, two rectified P-T values before the time of the detection are compared with Th3 to search for a beginning time earlier than the detection. Typically, Th1 = 2×s', Th2 = 1.5×s', and Th3 = 1×s'.

When the onset time has been found, an estimate of the direction of the first break is made. Then the quality of the first break and of the onset time is estimated. This is done by presenting the rectified amplitude of the first break, together with the two rectified amplitudes on either side of it, as a function of s' to form five signal-to-noise ratios (see Figure 9). Next, the period and maximum amplitude of the first four cycles are estimated. All of this information (quality numbers, time, amplitude, period) plus the current value of s' is output. There are two ways to use the quality numbers: (1) to make a statistical evaluation of the pick and (2) to approximately reconstruct the waveform of signal onset -- the last three values in the series can be viewed as a quantification of the first three amplitudes of the detected waveform. Thus, in viewing the output, an analyst has the traditional information plus enough information to form a concept (1) of the noise environment, (2) of the validity of the pick, and (3) of the appearance of the filtered signal.

Blandford et al. (1981) have added signals recorded at Pinedale, Wyoming, to artificial noise to form a digital test tape. They (1981) reported results on a number of different detectors that operated on these data. Using our normal detector configuration, with our simplest filter (F1, Figure 2), we operated on the tape and obtained essentially the same results as the SDAC on-line detector. The SDAC detector uses a 6-pole 1-2 Hz bandpass filter. Both of these detectors gave about the same results as two analysts who viewed an analog display of the digital data (Figure 17).

The current (May 1983) day-tape operation at ASL is facilitated if we perform an average of 30 or less detections/day/station. At some of the stations where cultural noise often is prominent (such as BCAO, GRFO) we have reduced the sensitivity of the detector to fulfill this need. At these stations (particularly GRFO) the performance is not up to that of the better stations. Emergent signals often seen at GRFO also contribute to its substandard performance.

Several design changes are being explored and tested to enhance the capabilities of the detector. These include (1) an algorithm to pick secondary phases, (2) a modification to pick smaller events and yet maintain essentially a zero false alarm rate, (3) an algorithm to identify cultural noise, and (4) an algorithm to identify large teleseisms.

ACKNOWLEDGMENTS

Thanks are due Habib Merghelani (Ministry of Petroleum and Mineral Resources, Jeddah, Kingdom of Saudi Arabia) and Howell Butler (USGS, Jeddah). Robert Reynolds (ASL) cooperated in implementing the algorithm and coded it on the SRO. Jon Peterson (ASL), Sharon Fletcher (Sandia National Laboratories), Alan Steppe (Glendale, California) and Robert Blandford (SDAC) made many helpful suggestions on improving the manuscript. Wendell Green (formerly at ASL) contributed valuable suggestions on the strategy for implementing the algorithm. Subroutine PONE is a modification of a routine by L. Gary Holcomb (ASL). Barbara Kimball (ASL) contributed technical assistance.

REFERENCES

Blandford, R.R., O. Racine, and R. Romine. Single channel seismic event detection, Report VSC-TR-81-8, 7 Oct 1981, Teledyne Geotech, Alexandria, VA, 150 p.

Peterson, J., H.M. Butler, L.G. Holcomb, and C.R. Hutt (1976). The seismic research observatory, Bull. Seism. Soc. Am. 66, 2049-2068.

Unitech (1974). Operation and maintenance manual: seismic research observatory data recording system, Unitech, Inc., Austin, TX, 121 p.